



# Proceedings of the 2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech)



# 25-26 November 2015 Port Elizabeth, South Africa

Copyright and Reprint Permission: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the percopy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923. For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Operations Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

Copyright © 2015 by IEEE.

ISBN: 978-1-4673-7449-1 IEEE Catalogue Number: CFP1548-USB

## **COPYRIGHT**



# 2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech)

ISBN: 978-1-4673-7449-1 IEEE Catalogue Number: CFP1548U-USB

Copyright and Reprint Permission: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through:

Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to:

IEEE Copyrights Manager IEEE Operations Center 445 Hoes Lane Piscataway NJ 08854

All rights reserved.

Copyright © 2015 by IEEE.

# **AFFILIATIONS**

Welcome to the 2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech)

# Incorporating

The 26th Annual Symposium of the Pattern Recognition Association of South Africa and The 8th Robotics and Mechatronics (RobMech) Conference of South Africa

# Organized by

Advanced Mechatronics Technology Centre (AMTC) Faculty of Engineering, the Built Environment and Information Technology Nelson Mandela Metropolitan University

## In Cooperation with

IEEE Robotics and Automation Society and the Council for Scientific and Industrial Research

http://robmechprasa2015.nmmu.ac.za/ 2015prasarobmech@gmail.com



### **MESSAGE FROM GENERAL CHAIRS**

It is with the greatest of pleasure that we welcome you to the 2015 PRASA-RobMech International Conference, hosted at the Nelson Mandela Metropolitan University in Port Elizabeth, South Africa. This conference represents the coming together of two of South Africa's premiere research events, by incorporating the 26th Annual Symposium of the Pattern Recognition Association of South Africa (PRASA), and the 8th Robotics and Mechatronics Conference of South Africa (Robmech).

PRASA has a long history as the flagship pattern recognition conference in South Africa, covering all aspects of artificial intelligence, machine learning, and data science in areas such as image, speech and text processing, and often applying these to uniquely African problems. Although it has not been running for anywhere near as long, RobMech has similarly been the go-to destination for local research in all aspects of robotics and automation, again often with particular applications in addressing local problems.

Uniting these two communities provides a great opportunity to enhance the strengths of both, and increase the impact and significance of their respective contributions. Furthermore, there is much overlap in many of the underlying principles and technologies. South Africa has many great research institutions producing high-impact work in multiple different areas. Between representatives from all over the country, and fresh insights from abroad, we really hope that this conference will provide a forum to host exciting new discussions, and grow connections between fields.

We would like to take this opportunity to thank the Deon Sabatta and Marelie Davel for their commitment and hard work as the Programme Committee, as well as Karl du Preez and Eunice Marx for all the support as the Local Organising Committee. We'd also like to thank Riaan Stopforth for all his valuable advice, and everyone else who has been instrumental in pulling this conference together.

On behalf of the entire organising committee, the Advanced Mechatronic Technology Centre at Nelson Mandela Metropolitan University, the South African chapter of the IEEE Robotics and Automation Society, the SAIMechE, and the SAIEE, we hope you have an exciting, productive and memorable time at PRASA-RobMech 2015.

### Theo van Niekerk and Benjamin Rosman

### **PRASA-RobMech 2015 General Chairs**

## **REVIEW PROCESS**

This year, the combined PRASA-RobMech conference attracted papers from 151 authors originating from 12 different countries.

Full papers included in the proceedings all passed a blind peer review process:

- Track chairs were assigned for the main themes of the conference (robotics, mechatronics, speech and language processing, image processing/vision and general machine learning).
- Track chairs matched submissions and reviewers, ensuring that at least two credible reviews were obtained per submission.
- Reviewers commented on and scored the paper with regard to technical quality, relevance, clarity of presentation and contribution to the field.
- Final decisions on the inclusion of the papers were made by the joint technical committee.
- Reviews were returned to the authors. For papers that were conditionally accepted, authors were required to address reviewers' comments prior to submitting their camera-ready papers for final publication.

The review process for work-in-progress papers addressed the same criteria but was less strict, while retaining a strong focus on relevance for the PRASA-RobMech audience.

After disqualifying incomplete papers, 78 papers were submitted for formal review. During the review process, 177 individual reviews were processed. Of the reviewed papers, 45 (58%) were selected as full papers to be included in the proceedings. An additional 15 papers were allocated to a 'work-in-progress' session, where authors could obtain feedback on early work.

A big 'thank you' to the track chairs and 66 reviewers, who assisted us in ensuring the quality of this conference.

On behalf of the 2015 PRASA-RobMech programme committee.

Programme Chair: Robotics and Mechatronics Deon Sabatta, University of Johannesburg, South Africa

Programme Chair: Pattern Recognition Marelie Davel: North-West University, South Africa

## ORGANIZATION

### **GENERAL CHAIRS**

Theo van Niekerk, Nelson Mandela Metropolitan University, South Africa Benjamin Rosman, Council for Scientific and Industrial Research, South Africa

### LOCAL ORGANIZING CHAIR

Karl du Preez, Nelson Mandela Metropolitan University, South Africa

### SECRETARIAT

Eunice Marx, Nelson Mandela Metropolitan University, South Africa

### **TECHNICAL COMMITTEE**

**PROGRAMME CHAIR: ROBOTICS AND MECHATRONICS** Deon Sabatta, University of Johannesburg, South Africa

# PROGRAMME CHAIR: PATTERN RECOGNITION

Marelie Davel: North-West University, South Africa

### **TRACK CHAIRS**

Fred Nicolls, University of Cape Town, South Africa Thomas Niesler, Stellenbosch University, South Africa Riaan Stopforth, University of KwaZulu-Natal, South Africa Celik Turgay, University of the Witwatersrand, South Africa

# **PROGRAMME REVIEWERS**

Aderemi Adewumi, University of KwaZulu-	Raphael Angulu, Masinde Muliro University
Natal, South Africa	of Science and Technology, Kenya
Valentine Azom, University of KwaZulu- Natal, South Africa	Etienne Barnard, North-West University, South Africa
Udo Becker, Ostfalia University, Germany	Tracy Booysen, University of Cape Town, South Africa
Peter Bosscha, CSIR, South Africa	Willie Brink, Stellenbosch University, South Africa
Chris Burger, CSIR Meraka Institute, South Africa	Michael Burke, CSIR, South Africa
Oya Celiktutan Dikici, Queen Mary University of London, Great Britain	Mariette Conning, CSIR, South Africa
Shaniel Davrajh, University of KwaZulu- Natal, South Africa	Gerhard de Jager, University of Cape Town, South Africa
Hendrik De Villiers, Stellenbosch University, South Africa	Jason de Villiers, CSIR, South Africa
Alta de Waal, CSIR, South Africa	Febe De Wet, CSIR Meraka Institute, South Africa
John Dickens, CSIR, South Africa	Bernardt Duvenhage, CSIR, South Africa
Herman Engelbrecht, Stellenbosch	Alexander Ferrein, Aachen University of
University, South Africa	Applied Sciences, Germany
Tim Gebbie, University of the	Mehrdad Ghaziasgar, University of the
Witwatersrand, South Africa	Western Cape, South Africa
Igor Gorlach, Nelson Mandela Metropolitan University, South Africa	Natasha Govender, CSIR, South Africa
Jeremy Green, University of	Gerhard Gruhler, Reutlingen University,
Johannesburg, South Africa	Germany
Dieter Hendricks, University of the	Matthew Howard, King's College London,
Witwatersrand, South Africa	Great Britain
Felix Huening, Aachen University of	Stephan Kallweit, Aachen University of
Applied Sciences, Germany	Applied Sciences, Germany
Richard Klein, University of the	Ashley Kleinhans, CSIR, South Africa
Witwatersrand, South Africa	
Neil Kleynhans, North-West University,	Ben Kotze, Central University of
South Africa	Technology, South Africa
Chioniso Kuchwa-Dube, CSIR, South	Huseyin Kusetogullari, Gediz University,
Africa	Turkey
Louwrens Labuschagne, Stellenbosch	Jacques Marais, Stellenbosch University,
University, South Africa	South Africa

Michael Mitchley, University of the	Francesco Mondada, EPFL, Switzaland
witwatersrand, South Africa	
Andre Nel, University of Johannesburg,	Yohan Noh, King's College London, Great
South Africa	Britain
Markus Raffel, Aerospace Center (DLR),	Atanda Raji, Cape Peninsula University of
Germany	Technology, South Africa
Boitumelo Ramatsetse, Tshwane	Pravesh Ranchod, University of the
University of Technology, South Africa	Witwatersrand, South Africa
Reinhold Scherer, Graz University of	Stefan Schiffer, Aachen University,
Technology, Austria	Germany
Matthew Sibanda, Cape Peninsula	Willem Smit, Stellenbosch University,
University of Technology, South Africa	South Africa
Willie Smit, Stellenbosch University, South	Hamit Soyel, Queen Mary University of
Africa	London, Great Britain
Jules-Raymond Tapamo, Univesity of	Louis Ten Bosch, Radboud Univ Nijmegen,
KwaZulu-Natal, South Africa	Netherlands
Ewald Van Der Westhuizen, University of	Charl van Heerden, North-West University,
Stellenbosch, South Africa	South Africa
David Van Leeuwen, Radboud University	Pieter Van Vuuren, North-West University,
Nijmegen, Netherlands	South Africa
Dirk Van Compernolle, KULeuven - ESAT,	Beatrice Vaneden, CSIR, South Africa
Belgium	
Daniël Van Niekerk, North-West University,	Herman Vermaak, Central University of
South Africa	Technology, South Africa
Robyn Verrinder, University of Cape Town,	Serestina Viriri, University of South Africa,
South Africa	South Africa
Daniel Withey, CSIR, South Africa	Riaan Wolhuter, Stellenbosch University,
	South Africa
Kamil Yurtkan, Cyprus International	
University, Turkey	

# ACKNOWLEDGEMENT OF SPONSORS



advanced Mechatronic technology Centre













for tomorrow Department of Mechatronics



### **SUPPORTING SOCIETIES**







## **KEYNOTE SPEAKERS**

### Data Science in General, and a Case Study on Intensive Care Unit Monitoring

Professor of Machine Learning, School of Informatics, University of Edinburgh

Member of the Institute for Adaptive and Neural Computation, School of Informatics



**Chris Williams** 

### Speaker Biography:

Chris Williams is Professor of Machine Learning in the School of Informatics, University of Edinburgh, and Director of the Centre for Doctoral Training in Data Science. He is interested in a wide range of theoretical and practical issues in machine learning, statistical pattern recognition, probabilistic graphical models and computer vision. This includes theoretical foundations, the development of new models and algorithms, and applications. His main areas of research are in visual object recognition and image understanding, models for understanding time-series, unsupervised learning, and Gaussian processes. He obtained his MSc (1990) and PhD (1994) at the University of Toronto, under the supervision of Geoff Hinton. He was a member of the Neural Computing Research Group at Aston University from 1994 to 1998, and has been at the University of Edinburgh since 1998. He was program co-chair of NIPS in 2009, and is on the editorial boards of the Journal of Machine Learning Research and Proceedings of the Royal Society A.

### Abstract:

I will start by discussing Data Science in general, and the training programme followed by students in the EPSRC Centre for Doctoral Training in Data Science at the University of Edinburgh. In the second half of the talk I will show how methods from probabilistic machine learning can be used to tackle the major problem of detecting, cleaning and removing artifact from vital signs data collected from the Neuro ICU at the Southern General Hospital, Glasgow. I will describe the Factorial Switching Linear Dynamical System (FSLDS) and the Discriminative Switching Linear Dynamical System (DSLDS) models for this task. Joint work with: Konstantinos Georgatzis, Chris Hawthorne, Partha Lal, Martin Shaw, Ian Piper.

Travel grant supported by:



### Motor Skill Learning: From Simple Skills to Table Tennis and Manipulation

Professor, Technische Universität Darmstadt

Research Group Leader, Max Planck Institute for Intelligent Systems

### Speaker Biography:

Jan Peters is a full professor (W3) for Intelligent Autonomous Systems at the Computer Science Department of the Technische Universitaet Darmstadt and



Jan Peters

at the same time a senior research scientist and group leader at the Max-Planck Institute for Intelligent Systems, where he heads the interdepartmental Robot Learning Group. Jan Peters has received the Dick Volz Best 2007 US PhD Thesis Runner-Up Award, the Robotics: Science & Systems - Early Career Spotlight, the INNS Young Investigator Award, and the IEEE Robotics & Automation Society's Early Career Award. In 2015, he was awarded an ERC Starting Grant. Jan Peters has studied Computer Science, Electrical, Mechanical and Control Engineering at TU Munich and FernUni Hagen in Germany, at the National University of Singapore (NUS) and the University of Southern California (USC). He has received four Master's degrees in these disciplines as well as a Computer Science PhD from USC.

### Abstract:

Autonomous robots that can assist humans in situations of daily life have been a long standing vision of robotics, artificial intelligence, and cognitive sciences. A first step towards this goal is to create robots that can learn tasks triggered by environmental context or higher level instruction. However, learning techniques have yet to live up to this promise as only few methods manage to scale to high-dimensional manipulator or humanoid robots. In this talk, we investigate a general framework suitable for learning motor skills in robotics which is based on the principles behind many analytical robotics approaches. It involves generating a representation of motor skills by parameterized motor primitive policies acting as building blocks of movement generation, and a learned task execution module that transforms these movements into motor commands. We discuss learning on three different levels of abstraction, i.e., learning for accurate control is needed to execute, learning of motor primitives is needed to acquire simple movements, and learning of the task-dependent "hyperparameters" of these motor primitives allows learning complex tasks. We discuss task-appropriate learning approaches for imitation learning, model learning and reinforcement learning for robots with many degrees of freedom. Empirical evaluations on a several robot systems illustrate the effectiveness and applicability to learning control on an anthropomorphic robot arm. These robot motor skills range from toy examples (e.g., paddling a ball, ball-in-a-cup) to playing robot table tennis against a human being and manipulation of various objects.

Travel grant supported by:



# LIST OF PAPERS

	TITLE	PAGE
1	Modelling of Boiler Fireside Control in Flownex Software Environment	1
2	Optimizing Search and Rescue Missions through a Cooperative Mobile Robot Network	7
3	Using the Earth Mover's Distance for Perceptually Meaningful Visual Saliency	12
4	Development of an Educational Tool to Teach Primary School Pupils Multiplication Tables	19
5	A study on the effect of different image centres on stereo triangulation accuracy	24
6	The effectiveness of combining the likelihood maps of different filters in improving detection of calcification objects	30
7	An Intelligent Locator Inventory to Accommodate Mass Customization in an Automated Flexible Fixture System	37
8	Using a Classifier to Track the Edge of a Conveyor Belt	42
9	Developing a Java based RFID application to automate student attendance monitoring	48
10	Normal Distribution Transform Graph-based Point Cloud Segmentation	54
11	Prototyping during the requirements elicitation process in the development of an underground unmanned aerial system	60
12	Single-labelled Music Genre Classification Using Content-Based Features	66
13	Self-Learning of Inverse Kinematics for Feedforward Control of Intracardiac Robotic Ablation Catheters	72
14	Sliding Mode Control for Electromagnetic Levitation System Based on Feedback Linearization	78
15	Navigation and locomotion of a low-cost Automated Guided Cart*	83
16	Creating a distortion characterisation dataset for visual band cameras using fiducial markers	89
17	Topic models for conference session assignment: Organising PRASA 2014(5)	94
18	Off-Road Vehicle Active Suspension Using Fluidic Muscle	100
19	Extracting South African safety and security incident patterns from social media	106
20	Using generalized maxout networks and phoneme mapping for low resource ASR– a case study on Flemish-Afrikaans	112
21	Synthetic triphones from trajectory-based feature distributions	118
22	Review of Standard Rotor Configurations for a Micro Unmanned Aerial System	123

23	High-level Rapid Prototyping of Graphical Models	130
24	Unsupervised acoustic model training: comparing South African	136
	English and isiZulu	100
25	Introducing XGL - a lexicalised probabilistic graphical lemmatiser for	142
	isiXhosa	
26	BEETLE - A Modular Electronics Family for Robotics	148
27	Predicting vowel substitution in code-switched speech	154
28	Integration of an Electrical Discharge Machining Module onto a	160
	Reconfigurable Machine Tool*	100
29	Text-based Language Identification of Multilingual Names	166
30	A Local Planner for Ackermann-Driven Vehicles in ROS SBPL	172
31	HMM Adaptation for child speech synthesis using ASR data	178
32	Autonomous Prediction of Performance-based Standards for Heavy	101
	Vehicles	104
33	Rapid Prototyping of Small Wind Turbine Blades Using Additive	100
	Manufacturing	109
34	Gait Adaptation of a Six Legged Walker to Enable Gripping	195
35	Semi-supervised Spectral Connectivity Projection Pursuit	201
36	Face and Iris biometrics person identification using hybrid fusion at	207
	feature and score-level	207
37	Design of a Virtual Simulator for Transesophageal Echocardiography	213
	Remote Control	213
38	Emotional Feedback for Service Robots using Hapto-Acoustic Interface	219
39	Control Considerations of the Low Cost Prosthetic Touch Hand	225
40	Contrasting Classifiers for Software-based OMR Responses	233
41	Switching Cascade Controllers Combined with a Feedforward	220
	Regulationfor an Aggregate Actuator in Automotive Applications	239
42	Arachnid-inspired Kinesthesia for Legged Robots	245
43	A straightforward method for calculating the voicing cut-off frequency	252
	for streaming HNM TTS	232
44	Validating a Reconfigurable Assembly System Utilizing Virtual	258
	Commissioning	

2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech) Port Elizabeth, South Africa, November 26-27, 2015

# Modelling of Boiler Fireside Control in Flownex Software Environment

Mr. Andre Kellerman<sup>1</sup>; Prof. I.A. Gorlach<sup>2</sup> Department of Mechatronic Engineering Nelson Mandela Metropolitan University Port Elizabeth, South Africa s214360253@live.nmmu.ac.za<sup>1</sup> Igor.Gorlach@nmmu.ac.za<sup>2</sup>

Abstract - With the influx of renewable generation technology and the constraints placed on the grid due to unplanned maintenance on aging power stations, coal fired stations designed for continuous base load production have to fulfil the role of frequency support. With the high inertia of the coal fired power stations, the dynamic operating philosophy of frequency support will impact negatively on the overall efficiency and life expectancy of these plants.

The solution proposed is to develop a computational model of the boiler controllers capable of optimising boiler fireside control loops and control philosophies. The model will enable engineers to test and optimize various control parameters. The model will also be a safe environment for testing transient scenarios, such as the loss of a draught group or a multiple mill trip.

Keywords - Flownex; Boiler Fireside Control; Engineering Simulator; Boiler Control Modelling.

#### I. INTRODUCTION

The need for boiler control optimization and ability to simulate transient events has increased with more frequent plant failure, Independent Power Providers (IPP's) and the generation variability introduced by renewable generation capacity on the grid [1]. The implementation of the control model will enable optimization and philosophy testing without requiring the unit to be offline. Contingency plans can also be derived from the model response to transient states, such as the loss of a Draft Group (DG).

A literature review has been conducted to investigate aspects relating to boiler control, with specific focus on boiler control philosophies and mechanical plant being controlled.

The boiler fireside control loops will be implemented using Flownex®, and will include boiler controllers such as Unit Coordinator, Load Controller, Combustion Controller, Mills' Controllers, Forced Draught Controller, Induced Draught Controller and Primary Air Controllers.

The boiler control model will be tested and its output (set points) will be compared to historic plant data. Different modes are available for boiler operation [2], and with Boiler-Follow mode being the normal operating mode frequency support stations [3], it will be used as the active mode for the computational model.

The testing strategy entails using historic values as inputs to the model, with Automatic Generation Control (AGC) being simulated using the load controller set point values. The results obtained from the model will be compared to historic plant data, and the accuracy will determine the degree of optimization possible.

#### **II. PLANT FIRESIDE PROCESS**

The schematic of the fireside of the plant is shown in Fig. 1, whit a description of the process that follows.



Fig. 1: Plant fireside process overview [4]

Coal is received from the relevant coal mines via overland conveyor systems from where it is stored on a coal stock pile on the power station's grounds.

From the stockpile, a stacker-reclaimer is used to distribute coal to mill bunkers and then to mill feeders, which supply boiler mills with coal according to the required flow rate obtained from the Distributed Control System (DCS). The coal fed to the mills is then crashed and pulverized to ensure an increased surface area for combustion. Preheated Primary Air (PA), supplied to the mills by the PA fans, transfers the pulverised coal particles to the boiler via a network of ducts. The coal/air mixture ultimately passes through a burner before it enters the boiler furnace.

Secondary Air (SA) is supplied to the burner in order to provide the required stoichiometric ratio [3] and to ensure that there is an excess of air in the furnace to prevent explosions and overheating of the furnace. SA is produced by the SA fans and it is transported via common SA ducts, after being heated by the air heater. The SA is supplied to the row of burners, which are enclosed in a wind box. The required flow rate of SA is achieved by controlling dampers and vanes of the wind box.

After combustion has taken place, the flue gas is extracted into the atmosphere through a stack with the aid of Induced Draft (ID) fans. Before exiting though the stack, the flue gas passes through a particulate collection system, which removes most of the fly-ash. Some power stations also have  $SO_2$ scrubbers which remove the sulphur from the flue gas before it is emitted into the atmosphere.

All these processes are achieved by means of the DCS, the control functions and modes of which are described below.

#### III. PLANT FIRESIDE CONTROL

#### A. Control Modes

The DCS responds automatically to load fluctuations or when a signal is received from National Control by increasing/decreasing the firing rate of the boiler as required. Depending of the boiler mode in operation, there are four main boiler control modes, namely: Boiler-Follow, Turbine-Follow, Manual and Coordinated control.

The Boiler-Follow mode is considered as a normal operation mode for boilers that need to respond quickly to frequency deviations due to the load variations. For the Boiler-Follow mode, the boiler response follows the turbine response, while trying to keep the steam pressure supplied to the turbine constant by throttling the turbine control valve. Therefore, it is also referred to as a constant pressure mode. The control logic of Boiler-Follow mode is shown in Fig. 2 [1].



Fig. 2: Boiler-Follow turbine control mode [1]

#### B. Control Loops

The process of steam generation in power stations is classified as a multivariable process. This term characterises processes in which two or more mutually coupled variables are to be controlled. In the boiler-turbine system, the outputs (dependant variables) are directly influenced by the state of the input variables (independent variables) [2]. All these variables are mutually coupled; meaning a change on an independent variable will have an effect on one or more of the dependant variables.

These variables form the inputs and outputs to the control loops, with conditioning of the variables within various controllers residing in these loops.

A hierarchical structured approach [1] is used for the DCS of the boiler control, as shown in Fig. 3.



Fig. 3: Hierarchical model of the boiler control [3]

Fig. 3 represents the relationships between the various controllers within the control levels. The green lines are classified as controller outputs with the red lines being feedback received from other controllers.

The unit coordinator generates the required set point for the boiler combustion based on the turbine load and passes the set point via the load controller to the combustion controller, which controls the combustion process [4].

The combustion controller controls the process via three controllers on the group control level, namely: furnace pressure, total air and fuel controllers. These controllers in their turn control the sub-group controllers, such as fans and mill controllers. At the drive control level, the control of actuators takes place. As can be seen from Fig. 3, feedback is sent to the combustion controller in order to obtain closed loop control of the boiler-turbine system or the unit.

In this research, only the boiler fireside was modelled with the focus on the combustion process efficiency.

#### C. Description of the Controllers

The unit coordinator is responsible for the decision making of how the unit will be operated in terms of control modes selected. The unit coordinator receives the required load signal either from an operator or from the AGC and generates the set point.

The load controller receives load the set point from the unit coordinator and compares it to the maximum capability which the unit can deliver as well as an allowable ramp up speed.

The combustion controller is responsible for safe and reliable combustion, forming set points for the master coal flow and total air.

The total air controller and FD fan controller receive total air set point from the combustion controller and calculate the vane position of the FD fan inlet. The Secondary Air controller controls the SA register vanes at either side of the individual rows on the windbox.

The Mill Feeder controller controls the speed at which the volumetric feeder delivers coal to the pulveriser, with the PA controller controlling the amount of PA required to convey coal into the furnace.

The ID Fan controller is responsible for maintaining a safe furnace pressure, ensuring that flames will not escape and protecting against the possibility of the boiler imploding.

#### IV. MODELLING

Theoretical and computational models of the boiler control were developed, which are described below.

#### A. Theoretical Modelling

A theoretical model was derived from boiler control logic and descriptions [5]. Protections have been omitted in the construction of the model, as alarms will indicate whenever the unit becomes unstable.

Fig. 4 is an example of the theoretical approach on the model Mill Feeder controller, with a description that follows.



#### Fig. 4: Mill Feeder Controller

The Mill Feeder controller receives its coal mass flow demand set point from the combustion controller via the Mills Bias controller. The combustion controller distributes the Mega Joule (MJ) per second requirement to the online mills by taking into account the Calorific Value (CV) of the coal.

The Mill Feeder controller then alters the set point by adding the set point received from the Mill Differential Pressure (DP) controller. The value of the Mill DP controller is calculated by taking into account the distribution of coal within the mills and the differential pressure loss it induces, to prevent the mill from choking. The coal flow demand value in kg/s is converted to a feeder speed set point by dividing it with a feeder factor. This feeder factor takes into account the feeder speed constant (32.01 rpm), volumetric constant ( $0.03165 \text{ m}^3/\text{rev}$ ) and a coal density constant ( $950 \text{ kg/m}^3$ ), which converts the kg/s coal flow set point to a percentage of the maximum Mill Feeder speed.

The feeder speed set point is then compared to the actual feeder speed before the speed deviation is passed to a Proportional and Integral (PI) controller. The output of the PI controller is passed to a Variable Speed Drive (VSD) that alters the feeder speed according to the demand from the controller.

The output signal from the Mill Feeder controller will drive the PA Vane controller, FD Fans controller and the SA controller, thus is a crucial element in boiler control.

The PA Vane controller receives the coal flow set point from the Mill Feeder and supplies PA according to a load line determined during optimization.

The FD Fans controller controls the amount of Total Air supplied to the furnace by adjusting the FD Fan Vane according to the total fuel set point.

The SA controller takes into account the total coal flow per mill, and adjusts the dampers in the windbox accordingly.

#### B. Computational Modelling

The computational model constructed was implemented using Flownex® Simulation Environment's Control Library [9]. A hierarchical approach was taken in implementing the boiler fireside controllers from the Unit Coordinator down to the control of individual drives.

The Mill Feeder model was implemented using Flownex® Compound Component library. This allows the user to easily create a complex custom component, define inputs and outputs. Using the Compound Component makes duplication of Controllers easy for Mills A to E, PA Controllers A to E and SA Controllers A to E, Filters, Operator set point Controllers (OSPC), lead and lag circuits.

The Mill Feeder controller receives the following inputs: Mill Feeder Load set point; feeder speed constant, OSPC, DP over the mill and the measured feeder speed.

The coal flow set point is received and kept in ranges between 3.63 kg/s and 7.5 kg/s, which are the minimum and maximum initial coal flow rate SP's per mill. This is then multiplied by an operator settable gain between 90 % and 110 %. The Mill DP demand signal is added to the biased demand set point, and is then converted to percentage value which will be used to drive the VSD through a PI controller.

The controller outputs to a PI controller which controls the VSD to a percentage of the maximum feeder speed. The coal flow rate is also displayed on the Human Machine Interface (HMI) for operator reference.

Fig. 5 shows the Mill Feeder controller implemented within the Flownex® environment, with Fig. 6 showing the internal components of the Mills Feeder controller compound component.



Fig. 5: Mill feeder controller implemented in Flownex®



Fig. 6: Mill Feeder controller compound component

In order to control process, a number of GUI were developed within Flownex®, which replicate actual HMI screens that are used at power stations controllers. In future, this will reduce the learning curve for operators by presenting the simulator in a familiar environment.

The Boiler Combustion HMI, shown in Fig. 7, allows the operator to view the Unit Load set point in MW, MJ/s fuel requirement, calculated CV of coal, the air/fuel ratio and master demand signal to the mills. The interactive HMI screens also allow the operator to control/bias the mills individually, and change the input to the Mill Feeder controller according to the requirements.



Fig. 7: Boiler Combustion HMI in Flownex®

For research and study purposes, an admin HMI was created, which allows the user to execute various scenarios. Process variable and parameters can be easily changed without the necessity of accessing the control model. A toggle switch of the admin HMI allows the user to switch between simulator mode and running the model on historic data inputs, as well as it allows the user to trip boiler auxiliaries and change internally set limits. These study scenarios will be described below.

#### V. RESULTS

By switching the simulator from "Historic Plant Data" to "Simulator mode", the model can be utilized as a fully functional engineering simulator. Two transient scenarios have been simulated, namely a Mill and DG trip.

#### A. Mill Trip

The trends shown in Fig. 9 to Fig. 12 display a transient scenario of a Mill trip at 15 s when the unit was running at full load Maximum Continues Rating (MCR) at 200 MW with a Unit Efficiency set at 35 %. The Mills where biased individually at 90 %, 95%, 100%, 105% 110% respectively from Mill A to E for illustrative purposes.

After one mill trips, the unit capability is automatically reduced by 25 %, which reduces the Unit Load set point and MJ/s energy demand by 25 % as shown in Fig. 9. The operator Load set point remains unchanged at 200 MW. The new MJ/s energy demand is sent to the Combustion controller which is used in the calculation of the Master Demand signal.

By taking the CV into account, the Boiler Load Demand Fuel is converted to a total coal flow requirement set point, deemed the Master Demand signal. The Master Demand signal is then evaluated against the number of mills online, which calculates the new coal flow requirement per mill, as shown in Fig. 10. The coal flow requirement per mill is used to drive the PA and SA controllers. Additional liquid fuel was activated at 15 s, decreasing the coal flow requirement. This function was deactivated at 25 s as the mill was brought back into service at 30 s. The PA controller's response to the change in coal corresponds with the coal flow requirement trend, as shown in Fig. 11. The required PA flow rate is calculated by referencing the coal flow set point to a Mill load line, with minimum flows of 11.2 kg/s and 12.6 kg/s, depending on the Mill inlet air temperature. The Total Air set point shown in Fig. 12 follows the total energy requirement set point trend, which is shown in Fig. 9, with a reduction in total air requirement at 10 s and a gradual increase at 30 s when the mill was brought back into service.

The simulations results appear empirically correct as they follow relevant direction and amplitude.



Fig. 9: Boiler load demand fuel (Mill trip)



Fig. 10: Coal flow set point (Mill trip)



Fig. 11: Primary Air set point (Mill trip)



Fig. 12: Total combustion air required set point (Mill trip)

#### B. DG Trip

The trends shown in Fig. 13 to Fig. 16 display a transient scenario of a single DG trip at 50 s when the unit was running at the same conditions as discussed in the previous section of a Mill trip, while omitting the FO activation.

After the DG trips, the unit capability is reduced by 50 %, which reduces the Unit Load set point and MJ/s energy demand by 50 % as shown in Fig. 13. The operator Load set point remains unchanged at 200 MW. The remainder of the process follows the same steps to calculate the required coal, PA and Total Air flow rates. The FD fan was brought back into service at 300 s.

The simulation results appear empirically correct as they follow relevant direction and amplitude.



Fig. 12: Boiler load demand fuel (DG trip)



Fig. 13: Coal flow set point (DG trip)



Fig. 14: Primary Air set point (DG trip)



Fig. 15: Total Air required set point (DG trip)

#### VI. CONCLUSION

The computational model is currently in testing phase. Testing will be conducted to determine the feasibility of continuing with optimization of boiler control. Datasets from power stations will be used and samples would be selected for various fireside variables. The variables from the selected samples will then be used as inputs to the model. This will be done by accessing the historic plant data and simulating a transient event by using all relevant process variables. This will include measurements, internally set limits and set points, which excludes the necessity for a fully functional thermofluid model for the control model validation. With the process being simulated with real world data, the response of the model will be correlated to that of the control system of the selected power station. If the correlation is within allowable limits, integration can commence and the system can be tested as a whole.

#### ACKNOWLEDGMENT

The authors would like to acknowledge the support provided by Dr. Craig Boesack and Mr. Sam van der Walt for their contribution to the project.

#### REFERENCES

- [1] C. Potgieter, Automatic process control Basic, Jhb: Eskom, 2013.
- [2] G. Gilman, Boiler Control Systems Engineering, ISA, 2005.
- [3] J. &. S. Kitto, Steam: Its generation and use, 41st Edition ed., The Babcock and Wilcox Company, 2006.
- [4] S. Dukelow, Control of boilers, ISA, 1991.
- [5] Heselton, Boiler Operators Handbook, Liburn, Georgia: Fairmont Press, 2004.
- [6] F. D. Mello, Boiler Dynamics and Controls (Course notes).
- [7] H. Kristinsson and S. Lang, "Boiler control improving efficiency of boiler systems," Lund University, Lund, 2010.
- [8] J. Feltas, Boiler effects on steam turbine response, New York: Power Tech. Inc., 2004.
- [9] M.Tech Industrial, Flownex DCS Library Manual, Potchefstroom: M.Tech Industrial, 2014.

# Optimizing Search and Rescue Missions through a Cooperative Mobile Robot Network

A joint research collaboration between South Africa and Argentina

N. Naidoo, G. Bright, R. Stopforth Mechatronics and Robotics Research Group (MR<sup>2</sup>G) University of KwaZulu-Natal (UKZN) Durban, South Africa <u>nic.naidoo@gmail.com, bright@ukzn.ac.za,</u> stopforth@ukzn.ac.za

*Abstract*— Robots assist rescuers in search and rescue operations. An area of research that can significantly benefit search and rescue operations, if integrated, is cooperative robotics, since a multi-robot team can cover a vast area more efficiently than a single robot system. The research discussed in this paper focuses on the need for cooperative search and rescue robots, the related research in literature, and the design of a control architecture that will support the requirements of a cooperation system for search and rescue applications.

#### *Keywords—Search and rescue robots; cooperation; middleware*

#### I. INTRODUCTION

Search and rescue operations in disaster areas can be greatly enhanced with the aid of semi-autonomous robots. South Africa has a poor accident record in its mines and this is a practical case among others, where semi-autonomous mobile robots can go into unstable and toxic areas in mine shafts to rescue, identify dangerous areas, and in turn save accident victims and rescuers [1]. The technologies associated with search and rescue robots will also have a direct benefit to people who are affected in low economic housing in South Africa, and Argentina, and in many other circumstances. Fires can have a devastating effect in low cost housing developments and the search for injured people can be greatly enhanced with technologically advanced semi-autonomous search and rescue robots. In addition, search and rescue operations in other disaster cases like, floods [2], mudslides [3], earthquakes, avalanches, fallen buildings, explosions, forest fire, retrieval and deactivation of bombs, etc. can be greatly enhanced with the aid of semi-autonomous robots [4].

Semi-autonomous robots are used by rescue personnel to assist with search and rescue operations. These robots aid in locating and retrieving injured or trapped civilians in disaster areas which are dangerous or may pose a threat to human rescuers. Rescue operations by semi-autonomous robots can be performed quicker, more efficiently and with less risk to rescue personnel. The advantage of using semi-autonomous robots in search and rescue operations, viz. in mines, is that they are not Z. F. Zelasco, J. Donayo Department of Mechanical Engineering University of Buenos Aires (UBA) Buenos Aires, Argentina jfzelasco@fi.uba.ar, juddonay@gmail.com

susceptible to gases and temperature changes. If damaged, there is no injury or loss of human life.

The objective of the joint science and technology collaboration research between South Africa and Argentina is to research, design, develop, assemble and test different systems to improve semi-autonomous robot functionalities for search and rescue operations. Mining disasters, earthquake ravaged areas, nuclear meltdown zones, collapsed buildings and other disaster situations that were mentioned are areas where search and rescue robots are needed to assist rescuers with locating injured people. The significance of the research that is researched, designed and implemented may allow the semi-autonomous mobile robot to perform search and rescue operations in hazardous environments. This robot will eliminate the need for human intervention in an environment that would cause casualties and in many cases, death. The recent collapse of the Twin Towers buildings in New York (USA) [1], earthquakes in Haiti and Nepal, tsunamis in East Asia, nuclear disaster in Japan, mining disasters in South America and fires in low cost housing developments (shanty towns), highlight the need for investigating the way to improve the ability of rescue robots worldwide.

The following objectives of the research collaboration are anticipated:

- A technologically advanced system for semiautonomous robots that use the latest research and developmental innovation to provide vehicles that can be used in disaster zones by South Africa, Argentina and other countries worldwide for search and rescue operations.
- Research, develop and implement innovative technologies that can be used on semi-autonomous robots to enhance the quality of life of people who reside and perform rescue operations in South Africa and Argentina.

An area of research that can significantly benefit search and rescue operations, if integrated, is cooperative robotics. Multiple robot teams have a great advantage over single robot systems since they have the resources to perform complex tasks much quicker. In a search and rescue operation, a multiple mobile robot system can cover a vast area (such as a forest) in an efficient manner by functioning as a cooperative team with a common goal. In order to achieve cooperative success, an Artificial Intelligence (AI) system must be inherent in each team member (a decentralised control approach) or the AI forms part of a centralised/hybrid control system.

The purpose of this paper is to investigate the use of cooperative robotics in optimizing search and rescue missions, as part of the joint research collaboration between South Africa and Argentina. The paper is structured as follows: Section 2 introduces the field of cooperative robotics, identifies the building blocks of local navigation and control in robots, and discusses the need for a middleware service, covering a literature review of current research in the area of cooperation in search and rescue robotic operations. The section concludes by discussing the middleware platform that will be used in this research. Section 3 will see the design of a robot control architecture for search and rescue applications that will be integrated with the middleware. Section 4 discusses further research and concludes the paper.

#### II. LITERATURE REVIEW

#### A. Cooperation in robotics

The first step in the process of achieving robot cooperation is to identify the control architecture required in the design of such a system. Depending on the design of the robotic network, solutions may form high levels of communication and synchronization between robots (strongly cooperative), or allow for periods of functional independence among robots (weakly cooperative). The method of cooperation system applied to a robotic team introduces the mention of some typical control architectures: centralized, decentralized, hierarchical, and hybrid systems [5].

- Centralised architectures control the entire team of robots from a single point. This approach is unreliable due to a single point of failure as well as the impracticality of maintaining a high communication frequency between robots for real-time control.
- Decentralised architectures are a common approach to most robot teams where each robot assumes control based on a knowledge base of their local situation. Unlike the centralised system, this type of control is highly robust; however, it may be a challenging task to maintain global control of the team due to the local behaviours of individual members.
- Hierarchical architectures resemble military systems where each robot oversees the control of a group of robots. It is better than centralised architectures in terms of reliability but is still susceptible to team failure due to the dependence on robots structured higher up in the hierarchy.
- Hybrid architectures is a combination of centralized and decentralized forms where local control and higher level plans are achieved, thus establishing a robust and potentially efficient control system. Hybrid

architectures are applied in many multi-robot applications [6].

Due to the exploration activity of search and rescue mobile robots, individual robots in the team must have periods of operational independence for local navigation and control, however, with the idea of cooperation and global path planning, each robot should be in communication with a control centre; thus the control scheme adopted for the research is a hybrid architecture.

#### B. Local navigation and control

Navigation is one of the most challenging competences required of a mobile robot; four building blocks of navigation are identified [7]:

- Perception: the robot must interpret the raw data extracted from the sensors to obtain meaningful data.
- Localization: the robot must identify its position in the environment.
- Cognition: the robot must decide how to achieve its goal position in the environment by planning a path.
- Motion Control: the robot must drive its motor outputs to achieve mobility in the planned direction.

In addition to these building blocks, the robot must be able to avoid obstacles en-route to its goal. Fig. 1 illustrates a structure of the components required in the design of a semiautonomous mobile robot system. With reference to a cooperative search and rescue robot, the localization, perception, and a part of the cognition module should be under a local, decentralized control system. The "mission commands" and "cognition path planning" blocks of the cognition module relate to the global mission, hence they comprise of the centralized control system.



Fig. 1. Control structure for autonomous mobile robots

The heterogeneities of individual robots in the team (such as data from various sensors) can also be used by other team members to enhance the search and rescue process, thus increase the victims' chances of survival. In order to achieve cooperative success, the distributed network of robots must utilise a software service, commonly known as the middleware layer.

#### C. Middleware for search and rescue robotics

The middleware can be seen as the glue that links everything in the robotic network together; it should be designed in a manner that will allow for the easy integration of robots to the network, especially in a search and rescue application where system configuration time is a critical factor. Another area of interest that can influence the time duration of successful rescue missions is the intelligent use of robotic resources to optimize the search process; the middleware must promote robot cooperation by providing software interfaces that will mask robot heterogeneities (from sensors and actuators) and simplify the system integration.

The following sub-sections discuss the functionalities of a few middleware platforms in literature that can be used for cooperative search and rescue robot applications.

#### 1) CINeMA

The Cooperative Intelligent Network Management Architecture (CINeMA) middleware was developed by researchers to ensure network connectivity and cooperative localization for search and rescue robots in disaster areas [8]. The motivation of the research stemmed from the problem of weak radio signal transmission due to the absorption and reflection of signals in collapsed buildings or underground areas. In addition to this problem, the localization of robots is affected since the central monitoring station does not receive the robot's GPS signal and thus cannot identify its position.

The purpose of CINeMA is to ensure reliable network connectivity between the robotic nodes by monitoring the RSI (Received Signal Strength Indicator) and LQI (Link Quality Indicator), and then cooperatively moves the robots within a safe region of communication through the use of EKF (Enhanced Kalman Filter) and k-NN algorithms based on the radio signal strength. CINeMA also supports the reestablishment of a robot's localized position by using a two way ranging (TWR) technique between the "lost" robot and the other localized robots [8].

#### 2) Miro

The Miro middleware [9] is designed using an objectorientated approach and adheres to the Common Object Request Broker Architecture (CORBA) standard, which is a framework for developing and maintaining distributed software systems. Miro was applied in the research of heterogeneous mobile robot cooperation in search and rescue missions, where the capabilities of different robots were used to divide the search space by sharing sensor information and/or local maps [10].

The research also involved the development of pyMiro, a Python binding for Miro that enables the fast prototyping of Python coded algorithms, eliminating the need for complex C++ code implementations. Miro is not well known in the research community, thus its library code base and user support is underdeveloped in comparison to middleware such as Player or the Robot Operating System (ROS).

#### 3) Player

The Player Project, developed by B. P. Gerkey et al. [11], is a robotic network server that provides a transparent interface for robot and sensor control. The project is supported by robotic researchers and is widely used all over the world [12]. Player can be installed on a computer connected to the robot and provides interfaces to sensors and actuators over an IP network. Client application programs can communicate with the robot's hardware using standard interfaces via Player over a TCP socket.

It is also possible for a client program, located anywhere in the network, to access any device [11]; thus a particular robot can view the environment by use of another robot's sensor. The use of remote access to sensors can be useful when a robot agent utilizes another robot's LRF (Laser Range Finder), for instance, to gain knowledge of the environment map for the purpose of localization and navigation; this is particularly useful for robotic search and rescue applications, where the search process can be optimized thereby increasing the victims' chances of survival.

Player has been implemented by the authors as a middleware for the cooperation of heterogeneous mobile robot platforms in manufacturing environments [13], however, it has not been considered as the choice of middleware for this research due to the growing popularity, functionality, and support of the ROS middleware.

#### 4) ROS

The ROS middleware consists of nodes, messages, topics and services and nodes communicate with each other in a peerto-peer (P2P) manner by publishing messages and subscribing to published messages. An initial event called the "naming service" is centralized and relies on a master node, shown in Fig. 2. The communication sequence between a publisher and subscriber begins with 1) the publisher node registers the topic (e.g. a laser scan) to the master node, also referred to as the naming server, and informs the master about the entry point of topic data; 2) the subscriber node inquires the master on how to access the particular topic; 3) the master responds by sending the subscriber the entry point data, such as the host address and port number; 4) the subscriber now directly communicates with the publisher (host) via TCP or UDP connections, requesting for topic data; 5) the publisher responds to the subscriber by sending the topic data stream (e.g. laser scan data).



Fig. 2. Communication in ROS [14]

ROS has been developed in a modular manner organized in packages, which can contain nodes, configuration files, libraries, datasets, or third-party software. The idea behind this form of software organization is so that functionality of software packages can be easily integrated into the ROS framework. There is a wide variety of ROS packages available for various robotic implementations such as sensor/actuator drivers, robot path planning and navigation, robot simulation, and others [15]. Among these packages is "cv bridge", a ROS package that is responsible for interfacing ROS with OpenCV. OpenCV is a free open source computer vision and machine learning software library which has more than 2500 optimized algorithms that can be used for object and facial recognition, classification of human actions, tracking moving objects, produce 3D point clouds from stereo clouds, and many other applications [16]. The use of computer vision is a part of the collaborative research, thus the integration of ROS and OpenCV is one of the reasons why ROS was selected as the middleware for this research; the other reasons are listed as follows [14]:

- ROS is free and open source.
- The availability of a large software library for device driver implementations.
- ROS is a distributed system, so software nodes and user applications can run on different machines and they each can communicate with each other.
- There are numerous commercial robots powered by ROS [17].
- ROS is beginning to be integrated into industry through the ROS industrial consortium [18], and is backed by well known industrial players such as Yaskawa, ABB, BMW, and Siemens amongst others.
- ROS installs on the popular Linux Ubuntu distribution, and is actively developed and updated.
- Applications can be programmed in various programming languages.
- Due to its popularity, a great deal of help is available through wiki-tutorials, community forums and research papers.

#### III. CONTROL ARCHITECTURE DESIGN

The type of control architecture chosen for the design of the robot cooperation network is a hybrid scheme. Each robot will have operational independent control for fast response time activities such as mechanical control, obstacle avoidance and local path planning, whilst the higher level control is dependent on a control centre, namely the "Navigation Control Center" (NCC) as shown in Fig. 3. The purpose of the NCC is to 1) receive and decode map data from each robot in the search and rescue team, 2) update its record of the global map of the environment, 3) run a cooperation algorithm to determine the required surveillance area by a particular robot in order to optimize the search and rescue process (i.e. cover a vast area in the minimum amount of time), 4) send the required direction to the remote robot via the radio communication link.

Each remote robot in the network will be powered by the ROS middleware and consists of a master and subscriber/publisher nodes as discussed in Fig. 2. The reason behind using multiple masters in the multi-robot system is to

increase the reliability of the system and eliminate the single point of failure. Also, individual masters mean that each robot has direct, real-time control over its sensors and actuators; this is a necessity in both single and multiple robot systems.



Fig. 3. Robot communication and control architecture

The ROS based robots consist of the following software modules:

- Perception: ROS software drivers will be used for data extraction from GPS (for robot localization) and stereoscopic camera (for computer vision) sensors. The choice of sensors used for the research is based on multiple design factors, the discussion of which is beyond the scope of this paper.
- Localization: this software module will be responsible for accurately calculating the position of the robot in the environment, as well as dynamically building a map of the environment that specifies vacant and obstacle related areas that have been discovered by the robot. These discoveries must be shared with the NCC to ensure an optimized search and rescue process, thus the reason for communication with the "Encode data" module.
- Cognition: three functions are pertinent to this module; 1) the electrical control of the robot's motor drives through the implementation of mechanical system models, 2) avoidance of obstacles in the environment while considering safety measures to prevent harm to humans, and damage to the environment and the robot, 3) plan a local path in the approximate direction specified by the NCC, the objective being to survey an area in the search of humans that may require help or rescue.
- OpenCV algorithm: this involves the use of the ROS package for OpenCV integration, thus allowing for the application of vision algorithms that will be used to identify humans in the environment. Once again, this

information must be shared with the NCC (since rescue personnel can be sent to the location), thus the reason for communication with the "Encode data" module.

• Encode data: the purpose of this software module is to gather, encode and send 1) map data and 2) data related to the probability of human recognition. A design requirement for the development of this module is the use of data extraction and compression techniques to minimize the size of the data packet that will be transmitted to the NCC; this will limit bandwidth and communication problems that can contribute to a poor performance of the search and rescue process.

A design factor that is worth mentioning is the possibility of radio connectivity problems between the NCC and the robots due to poor signal strengths that are influenced by the structure and geography of the environment. This problem is a non-trivial one to solve and can involve an in-depth research analysis, as discussed by the researchers of the CINeMA middleware [8]. The assumption of good and acceptable radio communication will thus be taken, ensuring a focus on the objectives of this research topic.

#### IV. CONCLUSION AND FURTHER RESEARCH

This paper introduced the need for semi-autonomous search and rescue robots in disaster areas as well as the benefits of a cooperative multi-robot team. The idea of robot cooperation revealed the reality of robotic heterogeneities in the network that must be masked by the middleware layer. A few relevant middleware platforms were discussed as per the literature survey, among them being ROS, a popular and highly integrated platform in the robotics community.

One of the key benefits of ROS that suits the collaborative research is the integration of the OpenCV software library. A discussion on the control architecture design of the research described the software modules (OpenCV included) that will be involved in the hybrid control scheme comprising of remote ROS controlled robots and a central Navigation Control Centre (NCC). The design considered the key elements required for optimized search and rescue operations through cooperation.

Further research in this topic will involve the development and testing of cooperation algorithms coupled with robot model tests in a ROS simulation environment. Another significant part of the research collaboration is the implementation of vision algorithms which is currently being developed.

#### ACKNOWLEDGMENT

The research presented in this paper is part of a joint collaboration between the University of KwaZulu-Natal, in South Africa, and the University of Buenos Aires, in Argentina. The authors wish to thank their respective governments for funding this research and for making the collaboration possible.

#### REFERENCES

- D. Greer, P. McKerrow, and J. Abrantes, "Robots in Urban Search and Rescue Operations", © ARAA, 2002 Australasian Conference on Robotics and Automation, November 2002.
- [2] J. Casper and R. Murphy, "Human-Robot Interaction during the Robot-Assisted Urban Search and Rescue Response at the World Trade Center", IEEE Transactions on Systems, Man and Cybernetics, Part B, Vol. 33, No. 3, June 2003.
- [3] R. Murphy and S. Stover, "Rescue Robots for Mudslides: A descriptive study of the 2005 La Conchita Mudslide Response", International Journal of Field Robotics, Wiley, 2008.
- [4] K. Kleiner, "Better robots could help save disaster victims", January 2006.
- [5] N. Naidoo, G. Bright, R. Stopforth, Z. F. Zelasco, and A. Abeledo, "Navigation and Control of Cooperative Mobile Robots using a Robotic Middleware Platform for Industrial Applications", Submitted to the South African Journal of Industrial Engineering, 2015.
- [6] L. E. Parker, "Multiple Mobile Robot Systems", Chapter 40 of Springer Handbook of Robotics. Springer-Verlag, Berlin Heidelburg, pp. 921-941, 2008.
- [7] R. Siegwart and I. R. Nourbakhsh, "Introduction to Autonomous Mobile Robots", MIT Press, Cambridge, Massachusetts, 2004.
- [8] S. Chouhan, D. Pandey, and Y. Chul Ho, "CINeMA: Cooperative Intelligent Network Management Architecture for Multi-Robot Rescue System in Disaster Areas", in Proceedings of the International Conference on Electrical, Electronics, Computer Science, and Mathematics Physical Education and Management (ICEECMPE), pp. 51-61, New Delhi, India, December 2014.
- [9] S. Enderle, H. Utz, S. Sablatng, S. Simon, G. Kraetzschmar, and G. Palm, "Miro: Middleware for autonomous mobile robots", in Proceedings of IFAC Conference on Telematics Applications in Automation and Robotics, 2001.
- [10] D. Kruger, I. Van Lil, N. Sunderhauf, R. Baumgartl, P. Protzel, "Using and extending the Miro middleware for autonomous mobile robots", in Proceedings of the international conference on Towards Autonomous Robotic Systems (TAROS 06), Survey, UK, September 2006.
- [11] B. Gerkey, R. Vaughan, and A. Howard, "The player/stage project: Tools for multirobot and distributed sensor systems", in Proceedings of the International Conference on Advanced Robotics (ICAR 2003), pp. 317-323, Coimbra, Portugal, 2003.
- [12] About the Player Project, <u>http://playerstage.sourceforge.net/</u> wiki/index.php/Main\_Page, last accessed 02 August 2015.
- [13] N. Naidoo, G. Bright, and R. Stopforth, "The Cooperation of Heterogeneous Mobile Robots in Manufacturing Environments using the Player Middleware", Submitted to the R&D Journal of the South African Institution of Mechnical Engineering, 2015.
- [14] N. Naidoo, G. Bright, and R. Stopforth, "A Cooperative Mobile Robot Network in ROS for Advanced Manufacturing Environments", Submitted to the International Conference on Competitive Manufacturing (COMA'16), 2015.
- [15] List of ROS packages for Indigo, <u>http://www.ros.org/browse/list.php</u>, last accessed 18 July 2015.
- [16] About OpenCV, <u>http://opencv.org/about.html</u>, last accessed 02 August 2015.
- [17] Robots using ROS, <u>http://wiki.ros.org/Robots</u>, last accessed 18 July 2015.
- [18] ROS-industrial, <u>http://rosindustrial.org/ric-americas/current-members/</u>, last accessed 18 July 2015.

# Using the Earth Mover's Distance for Perceptually Meaningful Visual Saliency

Logan Dunbar<sup>\*</sup> and Fred Nicolls<sup>†</sup> Department of Electrical Engineering University of Cape Town Cape Town Email: \*logan.dunbar@gmail.com, <sup>†</sup>fred.nicolls@uct.ac.za

*Abstract*—Visual saliency is one of the mechanisms that guide our visual attention, or where we look. This topic has seen a lot of research in recent years, starting with biologically-inspired models, followed by the information-theoretic and recently statisticalbased models. This paper looks at a state-of-the-art statistical model and studies the effects of using a cross-bin histogram distance instead of the usually employed bin-to-bin histogram distance on saliency detection. Specifically, the earth mover's distance (EMD) is used to replace the Bhattacharyya coefficient, and it is shown that the EMD provides a more continuous distance function for the high-dimensional colour histograms. This feature also facilitates the removal of an implicit center bias from the model in question.

Both these modifications lower a standard set of saliency scoring metrics, which is attributed to the noisy and haphazard nature of the eye-tracking data. An explicit selection task is identified as a possible avenue from which more robust and meaningful saliency scores can be developed.

#### I. INTRODUCTION

Through evolutionary pressures and competitive advantage our eyes have become a primary conduit for information about our environment. One of the problems our vision system has to overcome is that of information overload. There are approximately 92 million rods and 5 million cones in an average human eye which would take enormous processing power to update and monitor in real time [5]. Instead, the eye does something known as sparse coding, in that it finds the minimum amount of information that adequately allows the brain to reconstruct the scene which gives us the best chance of survival. It does this by preprocessing the information into roughly 10-12 channels of abstracted information such as edges, motion and large areas of uniform colour in the scene [15]. In addition to this preprocessing phase, we have developed an attentional mechanism, which, as William James put it [11], "...is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought... It implies withdrawal from some things in order to deal effectively with others." This attentional mechanism is called visual attention when applied to the visual system and it determines what part of the scene we attend to.

There are certain triggers that guide this visual attention, the most popular theory being the corresponding bottom-up and top-down attentional mechanisms [24, 13, 10, 22]. The theory suggests that we have two streams of attention, one being

a highly parallel feature and stimulus driven mechanism, the other being a much higher-level goal-driven mechanism. These two systems work in tandem to decide what part of the visual field should be allocated the majority of our limited processing power. Sharp edges, abrupt colour changes, repeating patterns, moving objects and even seemingly higher-level concepts like faces can trigger the bottom-up mechanism to indicate interesting content [12]. Top-down goals such as looking for words, or a specific colour or face in a crowd, would basically prime the bottom-up process to only respond to those cues, overriding its default behaviour. This process of detecting the interesting part of the scene is known as visual saliency.

In recent years, researchers have developed models for saliency which have moved from the biologically-plausible models [13, 18, 10, 9, 22] into the information-theoretic [4, 6], statistical [27, 17] and transform-based [7, 16] models, each providing new and unique understandings of the way visual saliency works. Biologically-based saliency is a great place to start, due to it being successfully implemented in nature and therefore readily available to be studied, but perhaps through biological limitations or evolutionary pressures the current systems are suboptimal. The statistical models have recently shown good performance at predicting where humans might look, which is why this paper focuses on exploring these models and attempts to discover the effects certain assumptions and design decisions have on the resulting saliency predictions.

In particular, Liu et al. [17] provide a framework for computing spatiotemporal saliency that produces state-of-theart results. The framework relies heavily on histogram metrics to compute the similarity of local patches of the scene with a global representation of the scene, as well as with each other. The distance metrics used are of the bin-to-bin variety which are known to suffer from edge effects and are unable to deal with histogram shifts. Histograms become sparser as the dimensionality of the data increases, meaning these effects are even more pronounced when dealing with colour histograms. This paper proposes using the earth mover's distance (EMD) [20] to account for these effects. The EMD is a cross-bin distance measure, which makes use of a ground distance between bins to provide a more meaningful distance between histograms. It is shown that using cross-bin distance on a publicly available dataset perceptually improves the saliency maps, but reduces the saliency scores obtained using standard scoring metrics. This disparity is attributed to the bin-to-bin distance not adequately suppressing the saliency values of the background region coupled with the scoring metrics favouring less accurate and more spread out saliency maps due to the haphazard nature of eye-tracking data. The paper also looks at removing the implicit center bias from the model which reduces scores but again improves the perceptual output of the saliency maps. A novel experiment is proposed which uses the more specific click selections of participants by asking them to select multiple objects that stand out, and even extended to "colour in" salient objects, perhaps with a stylus, so that size and shape can also be accurately modelled. Section II reviews related work in saliency research, Section III details the experiments run and the results and conclusions from said experiments are discussed in Sections IV and V respectively.

#### II. RELATED WORK

Visual saliency is a measure of how much a certain stimulus stands out from its surrounds, or in other words how much it "pops out". Koch and Ullman [13] proposed a theoretical foundation for the saliency map, which indicates the saliency of a location in the visual field, and how it might be implemented in the primate brain. The first computational implementation of a saliency map was developed by Niebur and Koch [18] and further extended by Itti et al. [10]. Their model consists of an input image processed into a Gaussian pyramid [2], from which they compute intensity, colour, orientation and temporal change features. These features are linearly combined into a saliency map, giving a greater weight to the temporal change feature based on perceptual observations made by the authors. They select the most salient location in the saliency map by means of a winner-take-all process. Interestingly, they also implement an inhibitory return signal, reasoning that the saliency map does not find the most salient region and then stop, but rather it finds the most salient region, allows processing of it, finds the next most salient, allows processing of it, and so on. To allow for this, the inhibitory signal coming from the winner-take-all process applies a transient Mexican hat, or difference of Gaussians, at the location where the signal originated. This has the dual effect of inhibiting the most salient region but also of slightly raising the saliency of nearby regions, which might prevent the attention jumping too rapidly or drastically around the scene. (For the interested reader, Itti and Koch [9] provide a more thorough review of the biologically plausible saliency maps and their computational counterparts.)

Not all saliency maps are biologically based. Studies of the statistics of natural images show that there is an invariance to scale in natural images [21]. The property is known as the 1/f law and states that the amplitude A(f) of the averaged Fourier spectrum of an ensemble of natural images obeys a distribution  $E\{A(f)\} \propto 1/f$ . Hou and Zhang [7] show that the analysis of 2277 natural images revealed local linearity in the log spectrum of the images, with each image containing a similar trend with some statistical singularities. They reason that if there is a similarity in the log spectra across a wide

variety of images, the information that deviates from these smooth curves is what should be attended to. They therefore apply a local averaged filter to the log spectrum to generate the smooth trend curve, and compute the spectral residual of the image as the difference between the local average spectrum and the image spectrum. The authors claim that the spectral residual contains the innovation of the image, and it is this innovation which is defined as being most salient.

Li et al. [16] explored this concept further and discovered that the spectral residual is of little significance, and show that by replacing it with random white noise with the same average value and maximum as the spectral residual they are able to achieve almost the same saliency map. They deduced that the spectral residual, which can be approximated by a horizontal plane, actually acts a high pass filter on the image. The amplitude spectrum of natural images always has higher amplitudes at lower frequencies, so when the amplitude spectrum is replaced by a horizontal plane it is, in effect, treating all frequencies as equal. By virtue of this, the lower frequencies are suppressed and the higher frequencies are enhanced. This is almost equivalent to a gradient enhancement operation, which is why it discovers small salient objects but will only highlight the edges of larger objects and of textured regions in an image. They then turn the saliency identification on its head, and choose to search for nonsalient regions based on the fact that salient objects come in many shapes and forms, and can be spread across the image, whereas the backgrounds and nonsalient regions are generally repeating or uniform, which they then suppress to highlight the salient objects. The authors show that a repeated pattern in a signal corresponds to a sharp peak in the amplitude spectrum in the Fourier domain. Convolving the amplitude spectrum with a Gaussian kernel effectively suppresses the periodic background and nonsalient regions, leaving behind the salient objects which they highlight with some post-processing. The size of the smoothing kernel affects the size of the detectably salient region, so they introduce a scale-space representation and use the concept of entropy to select the appropriate scale. To include more features they replace the Fourier transform by the hypercomplex Fourier transform, using the opponent colour channels as the quaternion values.

Bruce and Tsotsos [4] approach saliency from an information-theoretic standpoint, defining saliency in terms of the self-information of local patches of the image with respect to their surrounds. To the authors, saliency is synonymous with surprise, or the expected number of guesses it would take to predict the local patch based on its surroundings. To achieve this, a bank of filters is learned from a database of natural images using independent component analysis (ICA), forming a suitable basis of Gabor-like filters that correlate well with the V1 cortical cells found in the primate visual system. An estimate of the distribution of each basis coefficient is learned across the entire image via non-parametric density estimation. The probability of observing a local patch centered at any image location is then evaluated by independently considering the likelihood of each corresponding basis coefficient, with

the product of all likelihoods yielding the joint likelihood of the entire set of basis coefficients. Shannon's measure of selfinformation is used to translate the joint likelihood into the resulting saliency map.

The rest of this paper focuses on the next model of saliency by Liu et al. [17], which, in a similar vein to Bruce and Tsotsos [4] above, uses global and local features to compute a saliency map. In particular, a statistical approach is taken, almost akin to outlier detection, whereby the global nature of the image is characterized by colour and motion histograms as features, and then compared via distance functions with smaller, homogeneous, edge-preserving regions called superpixels [1, 26] to generate the resulting saliency map. The model is made up of a colour saliency map and a motion saliency map, and is adaptively fused to generate a spatio-temporal saliency map. This paper looks at mitigating the shortfalls of the currently employed bin-to-bin histogram distance by introducing the cross-bin EMD and studies the effects on some standard saliency scoring metrics.

#### III. METHODOLOGY

This experiment introduces the EMD to the model of saliency found in [17]. Specifically, the bin-to-bin Bhattacharyya coefficient, which provides an indication of similarity between histograms, is replaced by a similarity value based on the cross-bin EMD, which generates a more continuous distance value and takes advantage of the perceptual uniformity in the underlying CIE  $L^*a^*b^*$  (CIELAB) colour space. This paper studies the effects of the EMD on the colour saliency, which extends naturally to the temporal saliency calculation by virtue of using the same formulations. Section III-A reviews the colour saliency calculation, Section III-B details how the saliency models are scored and introduces the dataset used for the experiment, and Section III-C details the experiments performed.

#### A. Colour Saliency

Each frame or image  $F_t$  is transformed into the CIELAB colour space, and is segmented into superpixels  $sp_i(i = 1, ..., n)$ , where n is the total number of superpixels. Each of the frame's CIELAB channels is uniformly quantised into  $q_b$  bins, generating a colour quantisation table CQ with  $q_C = q_b \times q_b \times q_b$  bins, with  $q_b = 16$  as per the original paper. Using CQ, the frame-level colour histogram CH<sub>0</sub> is calculated using the entire frame's pixels, and normalised such that  $\sum_{k=1}^{q_C} CH_0(k) = 1$ . The quantised colour for each bin, qc(k), is calculated as the mean colour of all pixels that fall into bin k. Superpixel-level histograms,  $CH_i$  (i = 1, ..., n), are then calculated and normalised such that  $\forall sp : \sum_{k=1}^{q_C} CH_i(k) = 1$ .

Liu et al. [17] make two assumptions to generate their colour saliency map: 1) salient regions generally show contrast with the surrounding background regions, and 2) salient object colours are generally more sparsely distributed over the scene than background colours. They quantify the first assumption as the global contrast in the frame, which is defined by comparing each superpixel-level colour histogram with the frame-level histogram

$$S_{\rm GC}(sp_i) = \sum_{j=1}^{q_C} \left[ \operatorname{CH}_i(j) \sum_{k=1}^{q_C} \|qc(j) - qc(k)\|_2 .\operatorname{CH}_0(k) \right]$$
(1)

where  $\|\cdot\|_2$  is the  $L_2$  norm. This states that the global contrast for a superpixel in relation to the frame is calculated as a sum of occurrence-weighted distances between the quantised colours present in the superpixel and the frame.

To quantify the second assumption of the colours of salient objects being more sparsely distributed, they define the spatial sparsity measure. To compute the spatial sparsity, each superpixel is compared with every other superpixel to create an intra-frame similarity value

$$\lambda_{intra}(sp_i, sp_j) = \sum_{k=1}^{q_C} \sqrt{\operatorname{CH}_i(k).\operatorname{CH}_j(k)} \cdot \left[1 - \frac{\|\mu_i - \mu_j\|_2}{d}\right]$$
(2)

where  $\mu_i$  and  $\mu_j$  are the centroids of  $sp_i$  and  $sp_j$  respectively, and d is the diagonal length of the frame. The first term is the Bhattacharyya coefficient, which measures the similarity between the two superpixel colour histograms, and the second term is a distance weighting function. The equation will evaluate higher for superpixels with more similar colour distributions to one another, and which are spatially closer to one another. The rest of the saliency calculation is as per the original paper [17].

#### B. Scoring Saliency

The most common form of saliency model test is to use the free-viewing task [14], which is accomplished by tracking the eye movements of human subjects using commercial grade eye-tracking systems while they freely view image or video databases. More recently, specific task-based viewing and object segmentations have also been used. A number of measures for how well a saliency map predicts or accounts for spatial attention have been developed based on these eyetracking data.

In this paper, the correlation coefficient (CC), the normalised scanpath saliency (NSS) [19], the area under the receiver operating characteristic curve (AUC) [23] and the shuffled AUC (sAUC) [27] scoring metrics are used. To compute the CC, a heat map from the fixation data is generated by convolving the map of fixations with a Gaussian the same size as the high-quality foveal region, usually 2° of the visual field [8]. The CC is then computed between this heat map and the generated saliency map. Positive scores indicate correlation, negative scores indicate anti-correlation and 0 indicates no correlation. The NSS tests the correspondence of the human fixation points with the model-generated saliency maps. The model-generated saliency map is linearly normalised to have zero mean and unit variance, then the values of this normalised map at all fixation locations are averaged to provide the NSS score. Due to the normalisation, positive values indicate a greater than chance correspondence of the human fixations with the saliency map, zero indicates no correspondence and negative values indicate anti-correspondence. Receiver operating characteristic (ROC) is used to evaluate a binary classifier system by varying its discrimination threshold. The modelgenerated saliency map S is treated by a varying threshold on the saliency values, creating a binary fixation map for each level of the threshold. The human fixations are then used as the ground truth. The ROC curve is drawn as the false positive rate  $F_p$  (incorrectly labelling non-fixated locations as fixated) versus the true positive rate  $T_p$  (correctly labelling fixated locations as fixated), and the total area under the curve indicates how well the saliency model predicts human eye fixations. An AUC value of > 0.5 means the model is able to discriminate fixations from non-fixations greater than chance, 1 being perfect discrimination, and an AUC value of < 0.5means the model performs worse than chance, with an AUC value of 0.5 meaning the model contains no discrimination power at all. As a variation of this, the human fixations are taken as the positive set, and some uniformly sampled points from the image are chosen as the negative set [3]. A problem has been identified in the literature, common to all saliency evaluation methods, which has been termed the center bias [25]. It has been observed through many experiments that subjects' fixation points are biased toward the center of static images as well as in videos. One of the most prominent factors affecting the center bias is that of the photographer's bias, being that photographers generally place objects of interest towards the center of the frame. Another factor, known as the viewing strategy, is when subjects reorient upon new stimuli with greater frequency toward the center of the frame, usually after repeated exposure to photographer-biased stimuli. This is also due to many datasets requiring subjects to fixate on the center of the screen prior to being shown a new stimulus. The problem of center bias is made clear by Zhang et al. [27] and Judd et al. [12] when they use a centered Gaussian blob as the saliency map and obtain much greater than chance AUC scores, even higher than some saliency models. In the recent review of saliency scoring methods, Borji et al. [3] show that the center bias and smoothing of the fixation points into a heatmap affect all scores previously mentioned. They propose the sAUC as a viable measure, whose only difference is instead of taking a uniform sampling from the image as the negative set, all fixations from all other images are used as the negative set.

The image dataset<sup>1</sup> selected for this paper was introduced recently in [14]. It consists of 800 photos of both indoor and outdoor scenes, either taken by the authors or obtained from existing datasets and online search engines. The images were specifically chosen to contain lateral (left/right) contextual information of a tangible object.

What makes this dataset unique is that the authors provide eye-tracking data for three different tasks, as well as mouse click location data for an explicit selection task. The images were displayed to participants so as to subtend  $15^{\circ} \times 15^{\circ}$  of visual angle, and were shown on a grey background. First was the typical free-viewing task with recorded eye movements. Second, participants were asked to decide whether the left or right half of the image was more salient while tracking their eyes. Third, participants were asked to explicitly select the object or region from the images that they considered to be most salient using mouse click selections. And fourth, participants were cued with object descriptions prior to viewing an image, and asked to report whether the object was present or not, which was missing 50% of the time, also while having their eyes tracked. The definition of "salient" given to the participants was something that stood out or caught their eye. The authors gave an example of a red flower among a field of white daisies when prompted for clarification.

An interesting finding of this paper was that saliency models were better able to predict the explicit saliency judgement tasks. One of the possible reasons given is that free-viewing is not without a top-down goal, but rather each individual would have an intrinsic goal or set of goals in the absence of an extrinsic one, which could vary across participants, and even for the same participant over many trials. When told explicitly to determine the salient regions or objects, it is thought the goals of the top-down and bottom-up systems have aligned.

#### C. Experiment

The experiment is performed to determine the effects of using a more continuous, perceptually agreeable cross-bin distance in the saliency calculation, namely the EMD. The EMD is substituted for both the histogram distance formulae in (1) and (2), and the saliency calculated as per the rest of the paper. During the experiment, it was noted that in calculating the spatial spread of colour distributions around the frame an implicit center bias is introduced due to using the center of the frame as a reference point. This artificially inflates the scores due to it corresponding with the eye-tracking data. The more continuous nature of the EMD allows access to very similar colour distributions in the frame, based on perceptual distance in the CIELAB colour space rather than bin overlap as with the Bhattacharyya coefficient. A method was developed which, for every superpixel, would compute the EMD between its and every other superpixel's colour distribution, obtain a similarity measure by inverse normalising these distances, thresholding them to find very similar superpixels around the frame, and computing a sum of distances from their joint centroid:

$$D_{\text{EMD}}(sp_i, sp_j) = \text{EMD}(\text{CH}_i, \text{CH}_j),$$
(3)

$$\lambda_{\text{intra}_{\text{EMD}}}(sp_i, sp_j) = \frac{\max\left[\mathbf{D}_{\text{EMD}}(sp)\right] - \mathbf{D}_{\text{EMD}}(sp_i, sp_j)}{\max\left[\mathbf{D}_{\text{EMD}}(sp)\right] - \min\left[\mathbf{D}_{\text{EMD}}(sp)\right]},\tag{4}$$

$$SD(sp_i) = \sum_j \|\mu_j - \mu_c\|_2,$$
 (5)

where  $\mu_j$  is the centroid of a superpixel with  $\lambda_{\text{intra}_{\text{EMD}}}(sp_i, sp_j) \ge t$  and  $\mu_c$  is the mean of the centroids of all the superpixels above the threshold t. The final spatial sparsity is calculated as an inverse normalisation, as per the

<sup>&</sup>lt;sup>1</sup>The dataset can be found at https://labs.psych.ucsb.edu/eckstein/miguel/ research\_pages/saliencydata.html.

original paper. This is equivalent to switching the center of the frame in the original formula to the mean location of a superpixel's most similar superpixels. If the superpixel has similar superpixels spread around the frame, then  $\mu_c$ approaches the center of the frame and it works much like the original equation. However, if the superpixel has similar superpixels that are tightly bunched,  $\mu_c$  will be close to all of them and results in a small spatial distribution sum as expected. In the experiments t = 0.9 produced satisfactory results, and corresponds to patches being "90% similar in terms of colour" as defined by the EMD distance.

#### IV. RESULTS

Using the EMD drastically reduces the scores, similarly for removing the center bias, as seen in Figure 1. What is surprising is the increase in sAUC scores across all dataset tasks after removing the center bias. The sAUC measure was designed to cater for the center bias, which would naturally lower the original method's AUC scores, as is clearly seen. The large variance in the NSS scores actually points to an underlying issue with this saliency model, being that colour alone is not enough to capture saliency. When colour is the dominant salient feature in the image, then this model performs exceptionally well. However, when there are higherlevel features, such as faces or recognisable objects, then colour is no longer enough and the model performs poorly. Even though using the EMD reduces the scores almost wholesale, the saliency maps produced seem to match perception much better than the original method, with examples seen in Figure 2. Notice how the permeating grey background, induced by the bin-to-bin distance not being able to provide meaningful distances in higher dimensions, is reduced by using the EMD. This grey background is almost completely removed when the center bias is removed. By localising the spatial sparsity instead of using the center of the frame, it is better able to handle salient regions around the frame, which is only made possible by the continuous nature of the EMD.

Figure 2b shows fixations or click locations from each dataset task. The eye-tracking fixations are quite haphazard and spread around the image, which makes it difficult to accurately extract the salient regions in the image. By contrast, the explicitly clicked locations are very precise and cluster extremely well across participants. It is due to this fact that a novel scoring method based on the explicit click location is proposed. The explicit-click task should be extended to allow participants to select up to n salient regions in the scene, perhaps under some time limit. This aligns the bottomup and top-down goals to be searching for saliency, as well as providing the accuracy and precision to test pixel-level accurate saliency models. An extension of this test could be to allow the participants to "colour in" salient regions, perhaps on a tablet using a stylus of some sorts. This would give the added information of shape and size, and would further improve the scoring ability of the test.

#### V. CONCLUSIONS

Visual saliency is a crucial component to the way we as humans see. It is starting to be used in robotics applications to allow for autonomous navigation and continues to improve our understanding of the field of computer vision as a whole. This paper focuses on a state-of-the-art statistical-based saliency model and shows that using the cross-bin EMD preserves the desirable attributes of the underlying CIELAB colour space and produces more perceptually meaningful saliency maps, but lowers saliency scores. This also facilitates the removal of an implicit center bias which further improves the perceptual meaning of the saliency maps and again lowers scores. Based on the observation that explicit selection of salient regions provides perceptually meaningful, accurate and participant agnostic results, an extension to the explicit-click task is proposed which allows participants to select multiple salient regions, possibly under a time limit. A further extension would be to use a tablet or similar with a stylus to allow the user to select or "colour in" the salient regions, adding shape and size to the comparable saliency features.

#### References

- [1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels," Tech. Rep., 2010.
- [2] E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden, "Pyramid methods in image processing," *RCA Engineer*, vol. 29, no. 6, pp. 33–41, 1984.
- [3] A. Borji, H. R. Tavakoli, D. N. Sihite, and L. Itti, "Analysis of scores, datasets, and models in visual saliency prediction," in *IEEE International Conference* on Computer Vision, 2013, pp. 921–928.
- [4] N. D. Bruce and J. K. Tsotsos, "Saliency, attention, and visual search: An information theoretic approach," *Journal of Vision*, vol. 9, no. 3, p. 5, 2009.
- [5] C. A. Curcio, K. R. Sloan, R. E. Kalina, and A. E. Hendrickson, "Human photoreceptor topography," *Journal of Comparative Neurology*, vol. 292, no. 4, pp. 497–523, 1990.
- [6] D. Gao and N. Vasconcelos, "Decision-theoretic saliency: computational principles, biological plausibility, and implications for neurophysiology and psychophysics," *Neural Computation*, vol. 21, no. 1, pp. 239–271, 2009.
- [7] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8.
- [8] ISO 11664-1:2007, Colorimetry Part 1: CIE standard colorimetric observers. ISO, Geneva, Switzerland.
- [9] L. Itti and C. Koch, "Computational modelling of visual attention," *Nature Reviews Neuroscience*, vol. 2, no. 3, pp. 194–203, 2001.
- [10] L. Itti, C. Koch, and E. Niebur, "A model of saliencybased visual attention for rapid scene analysis," *IEEE Transactions on Pattern Analysis and Machine Intelli*gence, vol. 20, no. 11, pp. 1254–1259, 1998.



Figure 1. Saliency scores for the saliency model using the original and two newly proposed methods. Error bars indicate one standard deviation.

- [11] W. James, *The Principles of Psychology*. Henry Holt & Co., 1890, vol. 1, no. 2, ch. XI, pp. 403–404.
- [12] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *IEEE 12th International Conference on Computer Vision*, 2009, pp. 2106–2113.
- [13] C. Koch and S. Ullman, "Shifts in selective visual attention: towards the underlying neural circuitry." *Human Neurobiology*, vol. 4, no. 4, pp. 219–227, 1985.
- [14] K. Koehler, F. Guo, S. Zhang, and M. P. Eckstein, "What do saliency models predict?" *Journal of Vision*, vol. 14, no. 3, p. 14, 2014.
- [15] R. Kurzweil, How to create a mind: The secret of human thought revealed. Penguin, 2012, p. 95.
- [16] J. Li, M. D. Levine, X. An, X. Xu, and H. He, "Visual saliency based on scale-space analysis in the frequency domain," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 4, pp. 996–1010, 2013.
- [17] Z. Liu, X. Zhang, S. Luo, and O. Le Meur, "Superpixelbased spatiotemporal saliency detection," *IEEE Transactions on Circuits and Systems for Video Technology*, p. 1, 2014.
- [18] E. Niebur and C. Koch, "Control of selective visual attention: Modeling the "where" pathway," Advances in Neural Information Processing Systems, pp. 802–808, 1996.
- [19] R. J. Peters, A. Iyer, L. Itti, and C. Koch, "Components of bottom-up gaze allocation in natural images," *Vision Research*, vol. 45, no. 18, pp. 2397–2416, 2005.
- [20] Y. Rubner, "Perceptual metrics for image database navigation," Ph.D. dissertation, Stanford University, 1999.

- [21] D. L. Ruderman, "The statistics of natural images," *Network: Computation in Neural Systems*, vol. 5, no. 4, pp. 517–548, 1994.
- [22] D. D. Salvucci, "A model of eye movements and visual attention," in *Proceedings of the International Conference on Cognitive Modeling*, 2000, pp. 252–259.
- [23] B. W. Tatler, R. J. Baddeley, and I. D. Gilchrist, "Visual correlates of fixation selection: effects of scale and time," *Vision Research*, vol. 45, no. 5, pp. 643–659, 2005.
- [24] A. M. Treisman and G. Gelade, "A feature-integration theory of attention," *Cognitive Psychology*, vol. 12, no. 1, pp. 97–136, 1980.
- [25] P.-H. Tseng, R. Carmi, I. G. Cameron, D. P. Munoz, and L. Itti, "Quantifying center bias of observers in free viewing of dynamic natural scenes," *Journal of Vision*, vol. 9, no. 7, p. 4, 2009.
- [26] M. Van den Bergh, X. Boix, G. Roig, B. de Capitani, and L. Van Gool, "SEEDS: Superpixels extracted via energydriven sampling," in *European Conference on Computer Vision*. Springer, 2012, pp. 13–26.
- [27] L. Zhang, M. H. Tong, T. K. Marks, H. Shan, and G. W. Cottrell, "SUN: A Bayesian framework for saliency using natural statistics," *Journal of Vision*, vol. 8, no. 7, p. 32, 2008.



Figure 2. Example images with representative fixation overlays and their associated saliency maps generated using the original method and the two newly proposed methods.

# Development of an Educational Tool to Teach Primary School Pupils Multiplication Tables

Suzanne Smith<sup>1</sup>, Angus W.D.T. Steele<sup>1</sup>, Johan du Toit<sup>2</sup> and Mariette Conning<sup>1</sup>

Abstract—We present an educational tool that helps to teach primary school children their multiplication tables in a fun, interactive way. The design consists of a floor mat with numbers  $0 \cdot 9$ , to implement multiplication tables from the  $1 \times$  to the  $12 \times$ multiplication table. The mat is connected to a console which controls the system and displays the inputs from the mat. The numbers on the mat are activated by the children jumping on the numbers in sequence to enter a sequence of results for a given multiplication table. The first prototype, manufactured using low-cost and accessible components for both the mat and the console, proved to be functional and was successfully tested in a local primary school environment. The potential of this educational tool for improved mathematical learning for pupils from Grade 1 to Grade 4 is evident, and has sparked interest in further development of interactive educational mat-based tools.

#### I. INTRODUCTION

This work addresses the need for the development of skills and interest in mathematics, particularly for children in the early years of primary school where core mathematical knowledge is formed. The development and implementation of a fun educational tool using a mat-based approach provides an interactive system to teach children multiplication tables. The system was inspired by ideas from a teacher in the field and provides a solution to promoting mathematics for young learners.

Proficiency in mathematics is important in the understanding and development of science and technology, a strong goal for the South African government [1] [2], as well as for building the future on a global scale. An excellent review of the role of gaming in education is provided by McClarty et. al [3]. Further, mathematical game playing has been shown to be an effective method of enhancing student learning [4] [5], with research indicating that mathematical games improve understanding and success in maths [6] when taught in the classroom and practised at home in a fun and interactive manner. Maths tools, such as counters, blocks and abacuses have also been used with great success in teaching mathematics, as these tools assist students in applying creativity to mathematical problem solving [7]. The interactive mathematical tool presented in this work was designed with these aspects in mind and includes both interactive fun elements as well as abacus functionality to enhance the learning experience.

The use of gadgets for education has grown rapidly in the past decade, and allows for individualized learning, where students can work at their own pace using different resources [8]. Smartphones, laptops and tablets, interactive whiteboards, digital cameras and mp3 players are among the technologies employed to enhance learning through the use of gadgets [9]. More specifically, a recent survey of gadgets for E-learning mentions the following 5 gadgets for education: portable projectors, digital pens, tablets, digital microscopes, and fitness bands (for physical education and biology) [10]. Technology provides many positive effects on educational achievement, as described by Simuforosa [11], particularly where focus areas of the technology include inquiry (e.g. data modelling, spreadsheets), communication (e.g. word processing, simulations, tutorials), construction (e.g. robotics, computer aided design) and expression (e.g. interactive videos, music composition). Educational mathematical gadget-based games include products from Learning Resources<sup>®</sup> (same company that does Math Mat challenge) such as Minute Math Electronic Flash Card<sup>TM</sup> and the Math Whiz Maths Challenge<sup>®</sup> (products similar to our console).

A number of interactive and fun educational tools exist, including the Math Mat Challenge from Learning Resources<sup>®</sup>, which is most similar to the work presented here. However, many of these systems cater to a younger, pre-school audience and do not include multiplication tables. Mathematical foundations in the early years are of great importance. All children acquire some form of everyday mathematical capabilities [12] and this needs to be nurtured by providing opportunities for learning mathematics. By incorporating maths into play, this opportunity is heightened [12], highlighting the relevance of a fun mathematical teaching tool. The importance of incorporating a physical space to digital learning is described by Price and Rogers [13], where it is shown that combining physical activity with digital learning gives a child a more holistic learning experience.

The system presented in this work consists of an interactive mat with numbers which can be jumped on by students to form the correct sequence of numbers for the multiplication table of interest. The mat provides feedback in the form of colourful flashing lights and sounds, alerting the students when an answer is correct or incorrect. The teacher can control the operation of the mat by selecting the multiplication tables as desired and controlling the start and end of the game.

Initial implementation and testing of the system at New Horizons Private School - a local primary school outside of Groblersdal, located in the Limpopo province of South

 $<sup>\</sup>ast This$  work was supported by the Council for Scientific and Industrial Research, South Africa

<sup>&</sup>lt;sup>1</sup>Suzanne Smith, Angus Steele and Mariette Conning are with the Council for Scientific and Industrial Research, Materials Science and Manufacturing, Pretoria, South Africa mconning@csir.co.za

<sup>&</sup>lt;sup>2</sup>Johan du Toit is a teacher at New Horizons Private School, Groblersdal, South Africa piet@eagri.co.za

Africa - has received positive feedback from both students and teachers. Stemming from such use cases and invaluable inputs from users, a larger roll-out of the educational tool in primary schools is envisioned. Particular emphasis will be on primary schools in rural and under-resourced areas, where the impact on education has great potential through finding the fun in learning.

#### II. DESIGN

#### A. Functional system design

The main function of the educational tool is to help children learn their multiplication tables. This is achieved through a combination of sub-functions, which include a setup command, inputs for the children to provide their answers, and outputs to indicate to the children whether their answers are right or wrong. The functional design is shown in Fig. 1.



Fig. 1. Educational tool flow chart from a user perspective.

In terms of the functional design for the microprocessor, the first step would be to receive an input from the system to determine which multiplication table will be used. Next, the multiplication table is set up internally for real time comparison with inputs received. Inputs are then received from the mat as the child jumps on the numbers he or she wants to give as an answer. The inputs are compared with the multiplication table and depending on whether the answers are correct or not, a specific output is generated that will indicate to the child whether he or she was right or wrong. The child can try again until the correct answer is given. After the twelfth multiple is entered correctly, a number of lights flash in a sequence and a tune is played to indicate that the child completed the multiplication table successfully. This procedure can then be repeated numerous times as the child practises his or her multiplication tables. The logical flow is shown in Fig. 2.



Fig. 2. System functional diagram of the educational tool.

#### B. Technical design and components

Currently, the system consists of two separate parts, namely the mat and the console. The mat acts as a matrix of switches for inputs, and the console contains the main control program. The block diagram in Fig. 3 shows an overview of the system components for both the mat and the console.

An Arduino Mega2560 (Arduino<sup>TM</sup>, RS Components) is used for the microprocessor of the student console. The inputs from the mat are connected to respective input pins on the Arduino. According to the inputs received from the mat, the Arduino goes through a series of processing and gives outputs through the on board light emitting diodes (LEDs), speaker and seven segment displays. If the answer is correct, it gives a happy jingle, if the answer is wrong, it gives a sad jingle. The program steps through the entire multiplication table up to the twelfth multiple and if completed successfully, it plays a short, happy song with accompanying flashing LEDs.



Fig. 3. Block diagram showing overview of mat and console.

The mat is shown in Fig. 4 and consists of standard foam mat puzzle pieces that can be found at any local toy store. Conductive copper tape was placed between the top and bottom parts of the mat to make the switches. Each switch is connected to a specific input on the console so that each switch is associated with a specific button. One of the sides is connected to ground and the other one to the input on the student console. When a child jumps on a number, the input is pulled low and it registers as a button press.

The console shown in Fig. 4 includes a set of push buttons. These buttons can be used instead of the switches on the mat. The student can therefore use the console at his/her desk, as well as with the mat, where physical ability will be required to give answers.

As part of learning maths, an abacus is often used to teach the abstract concept of numbers. An abacus is included on the student console, represented by bright shining LEDs. These can be seen in Fig. 4. The green LEDs represent the units, the orange LEDs the tens and the red LEDs the hundreds. In Fig. 5 one can see the lights shining on the abacus depending on the number that was given as an input.

A teacher console was also designed and implemented, which can be used to control the student console remotely. A teacher would use this console to enter which multiplication table the student should complete. However, this module is not necessary for the system to function successfully, as the student console has the functionality to go into configuration



Fig. 4. Complete assembled system, consisting of the console for controlling the mat (top) and the interactive mat with numbers (bottom).

mode when "Y" (Yes) and "N" (No) are switched simultaneously. Due to time constraints, the full functionality of the teacher console was not completed before the field trials.

#### III. RESULTS AND IMPACT

The system design was manufactured using the various technical components described in section II. Once assembled and tested, the system was taken to a local school for teachers and primary school pupils to test the system and give inputs.

#### A. Console and mat functionality

Fig. 5 illustrates the functionality of the mat and console together to form the complete system. As an example, the  $3 \times$  multiplication table is used to illustrate a full sequence of operation, with shoes placed on the mat to represent the example selection performed by the student jumping on the mat.

Row 1 in Fig. 4 shows the configuration step to start a sequence for the mat. This is done by simultaneously jumping onto the "Y" and "N" blocks of the mat. This will trigger the configuration step, displayed as "Con" on the console display. In row 2 of Fig. 4, the multiplication table of choice is selected by jumping on the corresponding number, followed by the "Y" block to confirm the entry. The



Fig. 5. Example sequence of operation of mat and console for the  $3 \times$  multiplication table.

console displays a "3" to confirm the choice. Row 3 onwards shows the sequence of entering the  $3 \times$  multiplication table, starting with jumping on the "3", followed by jumping on the "Y" to confirm the answer. The console displays the number selected, and if correct, flashes the number on the display as well as the abacus LEDs and a short happy jingle is played through the speaker. In row 4 the next correct answer is entered by jumping on the "6" block. In the case of an incorrect answer, as illustrated in row 5 where an "8" is selected instead of a "9", the console displays three dashes and a sad jingle is played through the speaker. The child can try again to enter the correct answer (as in row 6, where a "9" is entered) and then continue with the sequence until the multiplication is complete (in this case up until  $3 \times 12 = 36$ ), as shown in row 7 of Fig. 4. When the final correct answer is entered by jumping on the "3" block, followed by the "6" block and then the "Y", the console flashes the answer, and all the LEDs flash repeatedly, with a happy tune played to show that the sequence was completed successfully.

#### B. Field testing and impact

The assembled system was tested by teachers and primary school pupils from Grade 1 to Grade 4 at the New Horizons Private School in Groblersdal. The mat proved to be robust and fully functional throughout a number of sessions. An example of school pupils using the mat and console together is shown in Fig. 6.



Fig. 6. School pupils using the mat and console at New Horizons Private School in Groblersdal and being supervised by teachers. Grade 1 (left) through to Grade 4 (right) pupils experimented with the mat.

The school pupils found the mat to be intuitive and fun, and quickly understood the concept of jumping on different numbers to create an ascending sequence of a selected multiplication table. The teachers were able to easily set the console and mat to a chosen multiplication table, showing the ease with which the educational tool could be adopted by teachers and school pupils.

#### **IV. CONCLUSION**

The educational mat and console provides an interactive tool for primary school children to learn their multiplication tables. This system can be used as a basis on which to develop a vast array of educational tools to assist children and teachers in day-to-day learning. The system could easily be expanded to include more complex mathematical functions.

Future work could include enabling the numbers and colours of the mat blocks to be shuffled, as children tend to learn patterns instead of the true answer, according to the teachers at the New Horizons school. This would imply that each of the numbers in the mat should have an individual sensor with pre-processing circuitry installed which will send a form of identification signal to the console. In this way the mat will be more modular and versatile.

The ability to store the results of each child using the system could allow for a high-score system to be implemented, giving children an incentive to improve their performance, while also allowing the teachers to track the progress of the learners.

This work highlights the advantages of incorporating fun into education, particularly for mathematics, as is evident from participation by primary school children.

#### ACKNOWLEDGEMENTS

The authors wish to acknowledge Tamsin Purkis for the mechanical design and manufacture of the console housing. The authors gratefully acknowledge the Council for Scientific and Industrial Research for providing funding for this work, as well as the New Horizons school in Groblersdal. The authors thank the school principal, Mrs Elmarie Fourie, and the learners for their time, inputs and participation in this project.

#### REFERENCES

- J. Butler-Adam. Education, training and innovation in the National Development Plan 2030. South African Journal of Science, 109(1/2), Art. a008, 2013.
- [2] Executive Summary National Development Plan 2030: Our future make it work, 15 August 2012.
- [3] K.L. McClarty, A. Orr, P.M. Frey, R.P. Dolan, V. Vassileva and A. McVay. A literature review of gaming in education, Gaming in Education, June 2012, pp 1-36.
- [4] C. Kamii and R. DeVries. Group Games in Early Education: Implications of Piaget's Theory. Washington, DC: National Association for the Education of Young Children, 1980.
- [5] M. Bakker, M van den Heuvel-Panhuizen and A. Robitzsch, Effects of playing mathematics computer games on primary school students' multiplicative reasoning ability, Contemporary Educational Phychology, vol. 40, 2015, pp. 55-71.
- [6] D. Holton, A. Ahmed, H. Williams and C. Hill, On the importance of mathematical play, International Journal of Mathematical Education in Science and Technology, vol. 32, 3(1), 2001, pp. 401-415.
- [7] J. Hiebert, T.P. Carpenter, E. Fennema, K. Fuson, D. Wearne, H. Murray, A. Olivier and P. Human. Making sense: Teaching and learning mathematics with understanding. Portsmouth, NH. Heinemann, 1997.
- [8] E. Mahajan. How gadgets are revolutionising education, The Times of India Tech, February 2012. Available online: http://timesofindia.indiatimes.com/tech/jobs/How-gadgets-arerevolutionising-education/articleshow/11911970.cms
- H. Gammuac. Classroom gadgets using technology to enhance learning. Calgary Herald, September 2013. Available online: http://calgaryherald.com/technology/classroom-gadgets-usingtechnology-to-enhance-learning
- [10] K. Patthamasoot and S. Charmonman. A survey of gadgets for Elearning. International Journal of the Computer, Internet and Management, vol 22 (3), 2014, pp 32-42.
- [11] M. Simuforosa. The impact of modern technology on the educational attainment of adolescents. International Journal of Education and Research, vol. 1(9), 2013.
- [12] H.P. Ginsburg, Mathematical play and playful mathematics: A guide for early education, 2006.
- [13] S. Price and Yvonne Rogers, Let's get physical: The learning benefits of interacting in digitally augmented physical spaces, Computers and Education, vol. 43, 2004, pp. 137 - 151.
# A study on the effect of different image centres on stereo triangulation accuracy.

Jason de Villiers Council for Scientific and Industrial Research, Pretoria, South Africa Email: jdvilliers@csir.co.za Robert Jermy Council for Scientific and Industrial Research, Pretoria, South Africa Email: rjermy@csir.co.za Fred Nicolls University of Cape Town Cape Town, South Africa Email: fred.nicolls@uct.ac.za

*Abstract*—This paper evaluates the effect of mixing the distortion centre, principal point and arithmetic image centre on the distortion correction, focal length determination and resulting real-world stereo vision triangulation. A robotic arm is used to generate a ground truth set of known positions resulting in 2078 measurements per cameras.

It is seen that compared to the naive use of the arithmetic image centre improvements of 10% to 27% in triangulation accuracy can be made by determining an optimal principal point. An optimal distortion centre has a smaller but still beneficial effect.

#### I. INTRODUCTION

This paper investigates the effect of using different 'image centres' on the accuracy of stereo vision triangulation. Different image centres have two effects on stereo vision: they can affect the accuracy with which the photogrammetric calibration can be achieved, and they alter the Three Dimensional (3D) vector created from the pixel position, focal length and pixel size. Wilson and Shafer [1] provide a complete taxonomy of image centres in a camera calibration system. The different image centres considered in this tudy are shown in Figure 1 and discussed below.

- 1) The arithmetic image centre: the point at half the image resolution horizontally and vertically.
- 2) The centre of distortion: the point in the image which is found to be optimal to use to correct the distortion using both radial and tangential corrections as described by Brown [2], [3].
- 3) The principal point: this is the point through which the optical axis (or chief ray) of the lens intersects the Charge Coupled Device (CCD).

In Figure 1 the spread of the different image centres and the radial and tangential distortions are magnified to highlight the effects. Note how the Two Dimensional (2D) vectors from the distortion centre and principal point vary markedly by several degrees. When turned into a 3D vector using focal length and pixel dimensions (refer to Figure 2) these vectors would result in extremely different points of closest intersection with another 3D vector. Note that in this study it was assumed that the CCD is orthogonal to the optical axis.

In order to test how these different image centre definitions affect the stereo accuracy, four cameras - of two different configurations - were used. Calibration data for each camera was captured so that its distortion parameters, principal point and focal length could be determined. Each camera also observed a set of points with known 3D positions from four different known 6 Degrees of Freedom (DOF) positions (hereafter referred to as a pose). The set of points observed from each pose consisted of three parallel planes roughly orthogonal to the cameras. Given that each camera observed these points from four known poses, there are 6 unique pairs of poses for each camera. Each of these pairs was used as a stereo baseline and a triangulated position for each of the known points was determined. A Root of the Mean Square (RMS) distance error for each camera was then calculated over all the known positions and pose pairs. This error is calculated for each of several different calibration variations where the distortion centre and principal point are either allowed to be optimized or fixed to either each other or the arithmetic centre of the image.

As stated earlier, the image centres affect not only the 3D vector created from an image point but also the photogrammetric calibration of a camera. The authors had previously shown [4] that the focal length calibration was both sensitive to noise and a major contributor in the accuracy of photogrammetric stitching of cameras calibrated with the proposed method [5]. In addition to quantifying the effect of 3D stereo triangulation accuracy, this work aims to determine if the improper usage of image centres is contributing factor of the focal length sensitivity.

The remainder of this paper is organised as follows: Section II defines the apparatus, calibrations and tests performed in the experiment. Section III provides and discusses the results of the experiment. Section IV places the results in context and presents the findings.

#### II. EXPERIMENT

This section elaborates on the apparatus and methods of the experiments performed. Section II-A expands on the photogrammetric camera calibration methods used. Section II-B defines the equipment and cameras used. Section II-C explains the accuracy measures resulting from the experiment.

#### A. Calibration overview

The calibration methods used in this work are based on the robot-arm based method of de Villiers and Cronje [5]. This meant that the following calibration sequence was performed:

1) Distortion calibration. The radial and tangential distortion coefficients and (optionally) the distortion



Fig. 1. Image Centres in 2D.



Fig. 2. Camera coordinate frame.

centre that best maps distorted image coordinates to their corresponding ideal pin-hole projection image coordinates were determined. Five radial and three tangential coefficients of Brown's model [2], [3] were used as this was shown [6] to yield the most accurate calibrations for real-world triangulations.

2) Focal length calibration. The method of de Villiers and Cronje [5] was extended to optionally determine

the best principal point by adding it as an additional two dimensions to be numerically optimized when determining the focal length that results in the most consistent set of camera poses calculated from observed tetrahedrons. The optimization was also enhanced from a simple linear brute force search to a non-linear optimiser, specifically the Leapfrog algorithm [7] was used. The centre used for this calibration was selectable to be either the arithmetic centre, or the distortion centre (if determined) or the optimal principal point.

3) Stereo triangulation accuracy. This calibration aims to determine the 6 DOF offset of the camera relative to its mechanical mounting interface. This is done by mounting the camera on a jig with four different non-aligned mechanical mounting interfaces at known poses and then observing a set non-planar points which are at known positions. The 6 DOF offset that minimises the global RMS triangulation error to the set of known non-planar points is then sought. The image centre used for triangulation can be any of those listed in Section I, provided that they were determined. It is this optimised triangulation error that is used as the error measure to compare calibrations performed with different image centres.

It is worth noting that the same calibration and triangulation data was used repeatedly for each camera with only different image centres used at each stage to determine this study. The same methods for calibration (more fully described in the patent [5] and the authors' previous calibration and related works [8]) were used for each calibration variation, thus removing this bias from the study.

#### B. Apparatus and physical configuration

The calibration apparatus used in this study is identical to that described by de Villiers and Cronje [5] and de Villiers et. al [9]. An ABB IRB 120 robotic arm with a stated 10 nm accuracy is used to present a light source in a repeatable sequence of known poses to the camera being calibrated/evaluated. A white Light Emitting Diode (LED) is used for visual cameras and a wire-wound resister observed through an aperture in a Teflon shield for the Long Wave Infrared (LWIR) cameras. The specifications of the cameras evaluated are given in Table I.

The distortion calibration of the camera used a robot movement sequence that completely subtended the Field of View (FOV) of the camera. An example of the recorded positions, that captured with Camera 1, is shown in Figure 3 where different columns are colour coded to emphasise the classical barrel distortion evident in the lens.

After the distortion calibration and lens focal length calibration the cameras observed a movement sequence of 363 points arranged in three parallel planes of 121 points in an 11 by 11 grid. The planes were a 100 mm by 100 mm square with a point captured every 10 mm horizontally and vertically. The three planes were separated by 25 mm. The captured image locations of the 3D grid of points from each of the four mounting positions is shown in Figure 4. The position of the mounting jig (using the first mount as a reference) and the relative positions of the other three mounts is given in Table II. The poses in Table II were determined with a Faro Arm Platinum portable contact measurement arm. The nominal distance from the cameras to the 3D grid of points was 400 mm.

TABLE I. CAMERA SPECIFICATIONS.

	Camera 1 & 2	Camera 3 & 4
Manufacturer	Xenics	Allied Vision
Model	Gobi 640 GigE	GT1920
Sensitivity Spectrum	8-14 μm	400-1000 nm
Resolution	$640 \times 480$	$1936 \times 1456$
Pixel size	17 μm	4.54 μm
Lens	Xenics 10mm	Pentax C814E
Nominal focal length	10 mm	8.0 mm



Fig. 3. Composite image of captured grid used for distortion calibration of camera 1.

#### C. Accuracy measure

The exact formulation of the stereo triangulation error for a specific camera and specific calibration thereof is given in Equation 1. The details on how to undistort an image point, turn it into a 3D vector in the camera Coordinate Frame (CF), use the jig, mount and camera relative poses to express the vector and camera pose in the robot CF and then triangulate the position of the light source in the robot's CF can be found in a good photogrammetry reference (e.g. [10]). Also note that the offset of the light source relative to the robot's end effector has to be taken into account or the orientation of the end effector has to be kept constant.

$$\mathbf{E}^{ST} = \sqrt{\frac{1}{N_C} \sum_{i=0}^{N_M - 2} \sum_{j=i+1}^{N_M - 1} \sum_{k=0}^{N_P - 1} \|T_{rp_k r} - T_{rp_k r}^{S,i,j}\|}$$
(1)

where:

- $\mathbf{E}^{ST}$  = the camera's stereo tringulation error,
- $N_C$  = the number of stereo triangulations per camera, =  $(N_M - 1)! \times N_P$ ,
- $N_M$  = the number of mounts on the jig (4 mounts),
- $N_P$  = the number of points in the robot movement sequence (363 points),
- $T_{rp_kr}$  = the known position of the light source at position k in the sequence, and

 $T_{rp_{rp}r}^{S,i,j}$  = the triangulated position of the light source at

	Yaw (deg)	Pitch (deg)	Roll (deg)	X (mm)	Y (mm)	Z (mm)
Mount 1	0.000	0.000	0.000	0.000	0.000	0.000
Mount 2	11.031	0.006	0.134	31.652	-170.365	-0.321
Mount 3	9.678	0.481	13.534	46.100	-117.429	-177.635
Mount 4	-4.662	3.735	-14.401	-70.239	-50.719	-150.796
Jig relative to robot	-169.322	-0.073	0.158	698.332	141.193	400.842

TABLE II. BRACKET AND JIG POSES.



(c) Bottom left mount image

(d) Bottom right mount image

Fig. 4. Composite images of captured points from the four locations using camera 1.

position k, using the camera on mount i and j respectively as a stereo pair.

#### D. Design of experiment

Three steps are required to determine the stereo accuracy of a given set of calibrations performed on a camera. These steps are the distortion calibration; focal length and principal point determination, and then actual triangulation. These steps must be performed in the sequence mentioned and are explained in more detail in Section II-A. There are also three image centres used (arithmetic, distortion and the principal point) and thus initially 27 different permutations of calibrations per camera. Many of the options are not possible because of the required order of the calibrations. Particularly, the principal point is only determined after the distortion calibration, thus the distortion calibration can only use either the arithmetic centre or the distortion centre (which is determined as part of the distortion characterisation). If the distortion calibration is forced to use the arithmetic centre then the distortion centre is never determined and thus the focal determination and triangulation steps cannot use the distortion centre. Similarly, if the focal length is forced to use either the arithmetic or distortion centre, the principal point is never determined and the triangulation cannot make use of it. This leaves the 10 possible calibration permutations, which are listed in Table VI.

Similarly, when considering the focal length determination there are only five possible permutations. These are listed in Table IV.

Distortion Calibration	Distortion (pixels RMS) Camera						
Туре	1 2 3 4						
None	5.511	3.947	2.431	2.539			
Arithmetic	0.271	0.112	0.514	0.239			
Centre							
Distortion	0.271	0.112	0.422	0.235			
Centre							

TABLE III. RESULTANT DISTORTIONS

#### III. RESULTS

This section presents the results of the experiment. The results of the two different distortions corrections on each of the four cameras is given in Table III. The results of the five unique focal length calibration is given in Table IV. Table V provides the determined image centres for each of the cameras. The RMS triangulation accuracy (see Section II-C) of the 4 cameras for the 10 variations of the image centres is given in Table VI.

The results for both calibrations for each camera as well as the baseline distortion measure are given in Table III. The values given are the RMS perpendicular distance over the image of the distance of each point from the best fit straight line through its row and its column. This is the same distortion measure that is minimised in the distortion calibration method (see Section II-A). Except for Camera 3, both distortion values yield similar residual distortions. This is in agreement with Stein [11] who stated that finding an optimal centre of distortion is a good approximation to finding the tangential coefficients. Although it has been shown [12] that there is still some benefit to be gained from finding both the distortion centre and tangential coefficient.

Table V shows how the different image centres changed. The reported principal point is the one calculated using a distortion calibration with its own optimised distortion centre. For the principal point the deviation from the arithmetic centre is given in three forms, the raw pixel distance, as a normalised percentage of the diagonal resolution, and angle in degrees. The angle is calculated using the pixel dimensions and nominal focal lengths given in Table I.

The achieved stereo triangulation accuracies are given in Table VI. It is worth recalling that the camera mount offset calibration [5] on which the stereo triangulation are based also

Centre used for		Focal Length (mm)					
Lens	Focal	Camera					
Distortion	Length	1	2	3	4		
Nominal Fo	cal Length	10.0	10.0	8.0	8.0		
Ι	Ι	10.009	9.539	7.666	7.703		
Ι	Р	10.000	9.830	7.755	8.387		
D	Ι	10.019	9.641	8.384	7.749		
D	D	10.034	9.798	8.384	7.552		
D	Р	10.009	9.844	8.118	8.337		

TABLE IV. DETERMINED FOCAL LENGTHS

D = Distortion Centre

I = Image Centre

P = Principal Point

refines the poses of the cameras and this explains why the large variations of up to  $1.5^{\circ}$  between image centres (and thus the vectors which are being triangulated) do not have a larger effect on the accuracies.

In all cases the naive approach of using the arithmetic centre is either the worst or near worst performing calibration. Determining an optimal principal point during the focal determination in all cases improved the resulting stereo vision accuracy and in Table IV typically resulted in focal lengths closer to the nominal value specified by the lens manufacturer.

When the principal has not been determined, the distortion centre typically yields slightly better stereo accuracies than using the arithmetic centre.

It was expected that determining the optimal distortion centre and principal point and then using the principal point for the triangulation would yield the best accuracy results. Inspection of the final two rows of Table VI shows that this is indeed the case for Cameras 1 and 3. Camera 2 achieved results only marginally better when instead using the arithmetic centre for the distortion calibration. The best accuracy for Camera 4 is achieved when using the distortion centre for the triangulation. These results are 6% better than those of the expected configuration.

#### IV. CONCLUSION

This study investigated the effects that mixing different image centres during camera calibration and operation had on the resulting real world triangulation accuracy with that camera. Four cameras of two different configurations (including two different sensitivity spectra) were used for the evaluation. For each camera, a common data set was captured for the calibration and evaluation of each camera. This data set was processed using 10 different permutation of image centres for the three steps required to calibrate and evaluate a camera. The evaluation of the camera's stereo triangulation accuracy compared the triangulated accuracy to the reported position of a robotic arm with a 10  $\mu$ m accuracy for a set of 363 points observed from four different positions. This yielded a total of 2178 real-world triangulation measurements per camera.

It was found that the naive case of using the arithmetic centre as the distortion centre and principal point is always among the least accurate methods. The error can be reduced by 10% - 27% by determining and using the optimal distortion centre and principal point. The majority of this improvement is attributable to the principal point. The optimal principal point also causes the calibrated focal lengths to be closer to the expected values.

#### REFERENCES

- R. Willson and S. Shafer, "What is the center of the image?" in Computer Vision and Pattern Recognition, 1993. Proceedings CVPR '93., 1993 IEEE Computer Society Conference on, Jun 1993, pp. 670– 671.
- [2] D. C. Brown, "Decentering distortion of lenses," *Photogrammetric Engineering*, vol. 7, pp. 444–462, 1966.
- [3] —, "Close range camera calibration," *Photogrammetric Engineering*, vol. 8, pp. 855–855, 1971.

Centre Type	Measure	Camera 1	Camera 2	Camera 3	Camera 4
Arithmetic	Coordinate	(320.000, 240.000)	(320.000, 240.000)	(968.000, 728.000)	(968.000, 728.000)
Distortion	Coordinate	(317.958, 238.470)	(312.423, 226.073)	(967.531, 723.960)	(958.032, 725.887)
Centre	Distance (pix)	2.552	15.855	4.067	10.189
	FOV %	0.638	3.964	0.336	0.841
	Angle (deg)	0.249	1.544	0.106	0.265
Principal	Coordinate	(319.358, 238.008)	(319.644, 228.388)	(951.547, 713.716)	(1000.65, 736.781)
Point	Distance (pix)	2.093	11.617	21.788	33.810
	FOV %	0.523	2.904	1.799	2.791
	Angle (deg)	0.204	1.131	0.567	0.879

TABLE V. IMAGE CENTRES

TABLE VI. STEREO TRIANGULATION ACCURACIES

0	Centre used for			Triangulation Accuracy (mm RMS)				
Lens	Focal	Stereo	Camera	Camera	Camera	Camera		
Distortion	Length	Triangulation	1	2	3	4		
Ι	Ι	Ι	0.502	1.500	2.104	1.736		
Ι	Р	Ι	0.496	1.195	1.989	1.242		
Ι	Р	Р	0.452	1.082	2.039	1.364		
D	Ι	Ι	0.512	1.376	1.679	1.666		
D	Ι	D	0.460	1.350	1.665	1.672		
D	D	Ι	0.523	1.208	1.671	1.985		
D	D	D	0.473	1.150	1.658	1.993		
D	Р	Ι	0.504	1.167	1.688	1.233		
D	Р	D	0.453	1.098	1.677	1.224		
D	Р	Р	0.452	1.087	1.604	1.329		
Best improvement (%)			9.96	27.87	23.76	29.49		
D-P-P impro	ovement (9	6)	9.96	27.53	23.76	23.44		

D = Distortion Centre

I = Image Centre

P = Principal Point

- [4] J. P. de Villiers and F. C. Nicolls, "A study on the sensitivity of photogrammetric camera calibration and stitching," in *Proceedings of* the 25th Annual Symposium of the Pattern Recognition Association of South Africa, ser. PRASA 2014, vol. 1, 2014.
- [5] J. P. de Villiers and J. Cronje, "A method of calibrating a camera and a system therefor," 11 2012, patent Number WO 2014083386 A2.
- [6] J. P. de Villiers, F. J. Wilson, and F. C. Nicolls, "The effect of lens distortion calibration patterns on the accuracy of monocular 3D measurements," in *Proceedings of the 22nd Annual Symposium of the Pattern Recognition Society of South Africa*, ser. PRASA2011, vol. 1, 2011, pp. 1–6.
- [7] J. A. Snyman, "An improved version of the original leap-frog dynamic method for unconstrained minimization: LFOP1(b)," *Applied Mathematics and Modelling*, vol. 7, pp. 216–218, 1983.
- [8] J. de Villiers, "Design and application of an automated system for camera photogrammetric calibration," Ph.D. dissertation, University of Cape Town, 2015.
- [9] J. P. de Villiers, R. S. Jermy, and F. C. Nicolls, "A versatile photogrammetric camera automatic calibration suite for multispectral fusion and optical helmet tracking," in *SPIE Defense, Security and Sensing*, vol. 90860W, 2014, pp. 1–9.
- [10] J. C. McGlone, *Manual of Photogrammetry, Sixth Edition*. Bethesda, Maryland: American Society for Photogrammetry and Remote Sensing, 2013.
- [11] G. Stein, "Lens distortion calibration using point correspondences," in Computer Vision and Pattern Recognition, 1997. Proceedings., 1997 IEEE Computer Society Conference on, Jun 1997, pp. 602–608.
- [12] J. P. de Villiers, F. W. Leuschner, and R. Geldenhuys, "Centi-pixel accurate real-time inverse distortion correction," in *Proceedings of the*

2008 International Symposium on Optomechatronic Technologies, ser. ISOT2008, vol. 7266, 2008, pp. 1–8.

# The effectiveness of combining the likelihood maps of different filters in improving detection of calcification objects

Keith Chikamai<sup>\*</sup>, Serestina Viriri<sup>\*</sup> and Jules-Raymond Tapamo<sup>†</sup> \*School of Mathematics, Statistics and Computer Science University of KwaZulu-Natal, Westville college Durban, 4000, South Africa <sup>†</sup>School of Engineering University of KwaZulu-Natal, Howard college Durban 4041, South Africa

Abstract-Breast cancer is the most prevalent form of cancer diagnosed in women. Mammograms offer the best option in detecting the disease early, which allows early treatment and by implication, a favorable prognosis. This study looks to combine the Wavelet, Median, Gaussian and a Finite Impulse Response filters for the task of detecting Malignant and Benign calcifications, which are among the primary indicators of breast cancer in digital Mammograms. These filters individually detect calcifications to varying degrees of success, but also create artifacts especially along the boundaries of curvilinear structures. They are combined in a way that improves overall detection, while diminishing their individual side effects. An Entropybased thresholding technique is finally used to segment the calcifications from the background. Experimental results show that the proposed model achieves a 100% detection rate, which shows the effectiveness of combining the likelihood maps from various filters in detecting calcification objects.

#### I. INTRODUCTION

Mammography is a popular approach in radiology that employs safe levels of X-ray radiation to highlight suspicious regions in the breast [3]. One of the common breast cancer indicators visible in a mammogram are calcifications; these are traces of calcium deposits in the breast ducts and lobules, which are visible in a mammogram as high intensity spots with spatial dimensions of between 0.05mm and 1mm. They generally fall into two categories: macrocalcifications, and microcalcifications. Macrocalcifications are relatively larger and rounder than microcalcifications; they are usually considered benign. Microcalcifications are of greater interest to researchers since they are highly indicative of breast cancer, especially if they occur in clusters [17].

Research activity with regards to microcalcifications is usually directed towards their detection, classification, enhancement, or a combination of these activities [16]. Wavelet analysis has received significant attention in the literature due to its versatility in characterizing microcalcifications. Rizzi et al. [15] implement a fully-automated wavelet-based Computer Aided Diagnosis (CAD) system based on the Biorthogonal (Bior 2.6) and Haar wavelets for the detection of microcalcifications on the MIAS database. They achieve a sensitivity score of 98% and specificity score of 89%. Jian et al. [10] implement the Dual-Tree Complex Wavelet Transform (DT-CWT) the classification and detection of calcifications, attaining a detection sensitivity of 83.5% and maximum classification score of 100%. Aquino et al [1] also exploit the DT-CWT (Real) for microcalcification detection attaining detection rates of between 20% and 66%. Detection of dense tissues proved the biggest challenge to the algorithm. Balakumaran and Shankar [20] employ a 1-dimensional coiflet transform with multi-scale analysis for the detection of microcalcifications, achieving an accuracy 96%.

In [21], the dyadic wavelet transform is combined with fuzzy shell clustering for microcalcification enhancement and microcalcification cluster detection respectively, achieving a detection rate of 95%. Challenges encountered included difficulty of detecting nodular-structured microcalcifications. In [5], the Haar wavelet coefficients are used to generate features for an extreme learning machine classifier towards detection of microcalcifications. Hamad et al. [8] investigate the optimal wavelet and decomposition levels for the detection of microcalcifications. The biorthogonal 2.2 wavelet is found to give the best positive predictive score at 82.6%. A later study [7] with similar objectives employs the biorthogonal 2.4 wavelet, improving the scores to 97.5% and 55.6% for the positive predictive value and specificity benchmarks respectively. In both cases, the authors propose further work to reduce false positive rates and improve sensitivity.

Common challenges in microcalcification detection in the literature presented above include undesirable properties of mammogram images such as, poor contrast, noise and aritificial objects. The density of the breast contributes in difficulty of microcalcification detection, with mammograms containing glandular and dense-glandular tissues proving relatively difficult to check for microcalcifications than those with Fatty tissues. This paper discusses the combination different feature maps as a means of increasing the detection sensitivity of microcalcifications in mammogram images. This model can find use in image processing algorithms requiring the detection of calcification-like objects as a preliminary step, and also in radiological applications for prompting potential suspicious areas for further investigation.

The rest of this paper is structured as follows: the proposed method is presented in section II. Section III discusses the results followed by the conclusion in section IV.

#### II. PROPOSED METHOD

Figure 1 presents the block diagram of the proposed system.



Fig. 1. Block diagram of the proposed method

#### A. Preprocessing

Mammogram images commonly contain undesired elements that might compromise the performance of algorithms if not addressed. Examples of these are artificial artifacts (like labels and markings), and noise introduced by instruments during image acquisition. The acquisition environment can also vary, implying differences in illumination that may affect contrastbased algorithms. In this work, the input image is preprocessed by simply subtracting from all pixels the mean intensity as follows,

$$I_{out} = I_{in} - mean(I_{in}) \tag{1}$$

where  $I_{out}$  is the output after the operation and  $I_{in}$  is the input (original) image.

#### B. Wavelet analysis

Wavelet transforms have been used for multi-resolution analysis of Microcalcifications in mammogram images. The transformations involve modeling the input signal as a superposition of wavelet basis functions, allowing the analysis of singularities and discontinuities in the signal. Unlike the Fourier transform which only allows the modeling of frequency information, wavelet transforms allow simultaneous signal space and frequency localization. This localization property allows for the isolation of noise, edges, or other discrete objects by filtering out the corresponding frequency. These properties make wavelets suitable for applications targeting localized high-frequency events or scale-variable processes. A formal treatment on wavelet analysis can be found in [5].

This work considers the Daubechies 1 (Db1) family for detection of microcalcifications; this is based on its relatively good performance in related works [5], [8], [7]. The following steps are taken during wavelet analysis,

- 1) Decompose the Input Image *I* by one level
- 2) Nullify the Approximation coefficients
- 3) Reconstruct the image using all the first level Coefficients to give the detail-enhanced image  $I_e$
- 4) Rescale the intensity range of the enhanced image to that of the original image to give the final image  $I_{wv}$

#### C. Gaussian/Median Filtering

The Median filter falls under order-statistics non-linear filtering approaches and is preferred for its preservation of edge information, which is lost in applications involving linear smoothing filters [6]. This makes it popular in applications that target denoising of images which contain salt and pepper noise [4]. Three square kernels are chosen to approximate the varying spatial extent of calcifications. The median filter response map is obtained using the following linear combination,

$$I_{Med} = I_{Med}(i,j) = \frac{2}{3}M_5(i,j) + \frac{1}{6}M_7(i,j) + \frac{1}{6}M_9(i,j)$$
(2)

 $M_d(i, j) = Med_{d,d}i, j$  denotes the square median filter with a spatial extent of  $d \times d$ . Calcification objects in the context of the images in this study are approximated at between  $3 \times 3$ and  $9 \times 9$  pixels, considering the digitization parameters of the MIAS database [19], and that their spatial dimension is reported to be between 0.05mm and 1mm [17]. This basis also informs the range of spatial dimensions used in the spatial-domain filters described in the following sections. From experimental runs, the larger spatial filters proved to amplify curvilinear structures at the expense of calcification objects, which led to segmentation challenges. Figure 2 exemplifies the Gaussian response maps for the smallest and largest kernels (the pattern is similar with the Median and FIR filters for corresponding kernel sizes). For this reason, the output of the larger dimension filters is weighted lesser than that of the smaller ones. The resultant image is subtracted from the original unfiltered image for isolation of Calcification-like objects.

The Gaussian filter was chosen for its similarity to the profile of Calcifications, which implies that it generates a strong response in the presence of calcification-like objects.



Fig. 2. Response map for the Gaussian filter for 2 kernels on case *mdb226*: (a)  $3 \times 3$  and (b)  $5 \times 5$ . Curvilinear structures are highly amplified in the second image. The images are combined to give (c)

The following gaussian equation is used to generate the kernel values, i.e.,

$$G_s(x,y) = e^{-\frac{x^2 + y^2}{2\sigma^2}}$$
(3)

A higher  $\sigma$  value increases the degree of attenuation. For the  $\sigma$  value, a set of 30 values linearly spaced between 2 and 5 were evaluated, with the best performance achieved at  $\sigma = 2.9$ . The filtering results are combined with the following weighting (using scalar multiplication) to give the Gaussianenhanced image  $I_{Gaus}(i, j)$ ,

$$I_{Gaus} = I_{Gaus}(i,j) = \frac{2}{3}G_5(i,j) + \frac{1}{6}G_7(i,j) + \frac{1}{6}G_9(i,j)$$
(4)

where,

$$G_d(i,j) = I(i,j) * Gaus_{d,d}$$
(5)

$$(i,j) \in S \text{ and } S \subseteq N^2$$
 (6)

 $Gaus_{d,d}$  denotes the Gaussian kernel with a kernel size of  $d \times d$ .

#### D. Finite Impulse Response (FIR) filter

This step uses three Laplace operators described in [22] to enhance microcalcifications.

-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	8	-1	-1	3	3	-1	-1	1	2	1	-1
-1	-1	-1	-1	3	3	-1	-1	2	4	2	-1
	(a)		-1	-1	-1	-1	-1	1	2	1	-1
				(b)		15	-1	-1	-1	-1	-1
							87 P	200	(c)		

Fig. 3. Finite Impulse Response filter kernels

The operators have dimensions  $3 \times 3$ ,  $4 \times 4$  and  $5 \times 5$ . The response map is obtained by filtering the input image with the FIR filters. The three resolutions were empirically chosen to cover the various ranges in size of calcification objects. The final filter response  $I_{FIR}(i, j)$  is obtained using the hadamard product of the individual FIR filters as follows,

$$I_{FIR} = I_{FIR}(i,j) = F_3(i,j) \circ F_4(i,j) \circ F_{5,3}(i,j)$$
(7)

where,

$$F_d = I(i,j) * FIR_{d,d} \tag{8}$$

$$F_{5,3} = F_5 - F_3 \tag{9}$$

$$(i,j) \in S \text{ and } S \subseteq N^2$$
 (10)

 $FIR_{d,d}$  is the FIR filter with the kernel dimensions  $d \times d$ . The filter kernel  $F_5$  amplifies curvilinear structures, which makes it difficult to segment calcifications. Equation (9) is thus used to diminish the strong curvilinear response.

#### E. Combination of filter responses

Each of the above filters provides a likelihood map for every pixel. Those pixels with high intensity values imply a high probability for presence of a microcalcification object. The image results from the convolution operations with the discussed filters are combined using the hadamard product to give the final calcification-enhanced image  $I_e$ , as follows,

$$I_e = I_{Med} \circ I_{Gaus} \circ I_{wv} \circ I_{FIR} \tag{11}$$



Fig. 4. Illustration of the combination of filter response maps for case *mdb233* (subimage (d)) showing the response maps for (a) Wavelet, (b) FIR, (c) Gaussian, (d) combined image as per Equation 11 and (f) the final thresholded image

The above operation significantly reduces smaller pixels values with an inverse effect on pixels having bigger intensity values. Furthermore, pixels flagged as suspicious by all the filters are significantly boosted while those that are not flagged by either of the operators are almost nullified. The simple scalar multiplication in Equation 11 was found to significantly reduce the effect of pixels that did not elicit a strong response from all the filters. This had the intended effect of nullifying artifacts that were unique only to a smaller subset of the filters. An illustration of this process can be seen in Figure 4.

#### F. Breast-Background artifact removal

All the filters used in this work provide a strong response along the boundary between the breast area and the image background in some of the images, which is an undesired side effect (An example can be seen in Figure 4a on the right side). This boundary artifact is eliminated by thresholding the filtered image using the mean value, followed by binary erosion (using a circular structuring element) of the image foreground as shown in Algorithm 1,

Algorithm 1 Remove Breast Area/Image Background Boundary Artifacts

1: **function** REMOVEBBARTIFACT( $I_{orig}, I_{proc}$ )  $t = mean(I_{orig}) \triangleright$  The threshold is the global mean 2: s = GetDiskStructingElement(5)3: ⊳ Disk structuring element of dimension  $5 \times 5$  for erosion  $mask = Thresh(I_{orig}, t) \triangleright Threshold I_{orig}$  using t 4: 5: mask = Erode(mask, s)▷ Erode Binary image using structuring element  $I_e = mask * I_{proc} \triangleright$  Enhanced image is a product of 6: the mask and the processed image 7: return *I<sub>e</sub>* 

8: end function

#### G. Removal of Small objects and Linear structures

Region size and eccentricity are used as criteria for the removal of small objects as well as linear structures. Objects having an area A < 3, as well as those having an eccentricity value E > 0.98 are removed as shown in Algorithm 2,

Algorithm 2 Removal of Small objects and Linear Structures
<b>Require:</b> $I(x,y) \in [0,1]$ $\triangleright I$ is assumed to be a binary
image
1: function REMOVEOBJECTS(I)
2: $A_{thresh} = 3$
3: $E_{thresh} = 0.98$
4: $Regions = GetRegions(I) \triangleright Regions$ represents
all image objects
5: for all $r \in Regions$ do
6: $A_r = GetArea(r) \triangleright$ Retrieve the object's area
value
7: $E_r = GetEccentricity(r) \triangleright$ Retrieve the object's
eccentricity value
8: <b>if</b> $A_r < A_{thresh}$ <b>or</b> $E_r < E_{thresh}$ <b>then</b>
9: $I(r(:)) \leftarrow 0 \qquad \triangleright$ Nullify pixels of region r
10: <b>end if</b>
11: end for
12: return I
13: end function

An eccentricity value of 0 represents a spherical shape with 1 being a linear object. Some noise objects were found to be near spherical, with high eccentricity values. Others were noted to have small areas of between one and two pixels. These two factors guided the choice of the eccentricity and minimum area thresholds.

#### H. Thresholding

Entropy information is used to determine the optimal threshold value for segmenting microcalcification objects from the background. This work customizes the Tsallis entropy thresholding technique discussed in [18] over a two-class problem, which is the maximization of the following function,

$$f(t) = Argmax \left[ S_q^A(t) + S_q^B(t) + (1-q) * S_q^A(t) S_q^B(t) \right]$$
(12)  
$$S_q^A(t) = \frac{1 - \sum_{i=0}^{t-1} \left(\frac{P_i}{P^A}\right)^q}{q-1}, P^A = \sum_{i=0}^{t-1} P_i$$
$$S_q^B(t) = \frac{1 - \sum_{i=t}^{L-1} \left(\frac{P_i}{P^B}\right)^q}{q-1}, P^B = \sum_{i=t}^{L-1} P_i$$

where, q is the entropy index, S the measure of entropy, L-1 the maximum gray level in the image region and  $P_i$  the probability of gray level i. This work considers segmentation of two classes: The breast background (A) and microcalcifications (B). The image background is ignored, which means that the search begins with t > 0, with t taking on the value of the minimum non-background pixel value. Before threshold calculation, the image is transformed to the range  $I(x, y) \in [0, 255].$ 

#### **III. RESULTS AND DISCUSSION**

#### A. Image dataset

The images used in this work were sourced from the Mammographic Image Analyis Society (MIAS) database [19]. The MIAS database has 29 ROIs containing microcalcifications, of which 15 are Malignant. These ROIs are spread over 25 different cases. This work extracted 27 ROIs classified positive for presence of microcalcifications, and 99 (randomly chosen) classified as normal, from the MIAS database each having a resolution of  $256 \times 256$  pixels with the cluster containing the abnormality centered on the image. The 27 ROIs were chosen from the database because they are clearly described in the ground truth with regards to center of abnormality and cluster size; this information lacks in the 2 ROIs that were left out.

#### B. Performance Metric

Sensitivity and specificity are common metrics for algorithms involving detection problems, especially in the medical domain [2]. These metrics, as well as other related ones, are based on True and False Positives (TP/FP). True or False refers to how the decision of the algorithm coincides with the true clinical decision. We base our measurement metric on Karssemeijer's criteria for counting True and False Positives [9]: A True cluster is flagged if three or more objects are detected within a radius of 1cm. A False Positive (FP) cluster is counted if none of the objects found in the cluster are inside the truth circle. The truth circle is determined from the ground truth provided together with the MIAS database. Having established the TPs and FPs, we calculate Sensitivity (also called the TP rate [11]) and Specificity as follows,

$$Sensitivity = \frac{TPs}{TPs + FNs}$$
(13)

$$Specificity = \frac{TNs}{TNs + FPs} \tag{14}$$

TABLE I SENSITIVITY AND SPECIFICITY RESULTS FOR MALIGNANT AND NORMAL IMAGES

	Sensitivity	Specificity
Normal/Abnormal ROIs	100	11
Malignant/Benign ROIs	57	40

#### C. Results

The system was developed using MATLAB R2014a and the performance benchmarks done on an Intel dual core i3 processor. In this section the results of the proposed model are presented. The white circle in the Figure 5-6 is an overlay delineating the cluster boundary as traced by an radiologist expert, according to the accompanying ground truth. Figure 5-6 illustrates three of the scenarios presented in the results. In Figure 5, all the microcalcifications in the clusters in the truth circle have been detected. The type of the breast tissue is Fatty-Glandular. In Figure 6, all the microcalcifications in the cluster have been detected. However, the proposed model also marked out other objects (lower right of the image) as calcifications.



Fig. 5. Database case *mdb209* - Malignant ROI with all microcalcifications in the cluster detected. (a) represents the original image. The (b) represents the enhanced image after linearly combining the output of wavelet analysis, Median and Gaussian filtering. (c) represents the final image, after thresholding and applying all the post-processing operations.

Table I presents the performance of the proposed model with regards to Sensitivity and Specificity measurements. The first row assesses the performance of the proposed model in detecting all microcalcifications, Malignant or benign, in a database containing 99 images of all pathology (Malignant+Benign+Normal). The second row presents the results of the proposed model's ability to detect malignant clusters in a database containing 27 malignant and benign cases.

#### D. Discussion

The sensitivity result in Table I that the proposed model positively identifies all malignant and benign calcifications



Fig. 6. Database case *mdb231* - Malignant ROI with microcalcification cluster detected as well as some False Positives. (a) represents the original image. The (b) represents the enhanced image after linearly combining the output of wavelet analysis, Median and Gaussian filtering. (c) represents the final image, after thresholding and applying all the post-processing operations.

according to the ground truth. The proposed model also scores above average in the discrimination of Malignant and benign microcalcifications. The major shortfall is that it falsely flags certain normal breast structures as calcifications, which gives the low specificity rate. The specificity rate can be understood under the context that this work set out to maximize the detection rate of microcalcifications where they exist. It should be pointed out that in most images where microcalcifications are present, they are uniquely detected by the proposed model without falsely flagging other parts of the same image as positive.

As attested to by the sensitivity results of the proposed model (Table I), all microcalcification clusters are detected and stand out from the background to a significant degree. In some cases, the proposed model falsely identifies noncalcification objects as microcalcifications, which is the cause of the low specificity rate. This point is best illustrated in Figure 6, where the microcalcification cluster is clearly distinguishable from the background after the filtering stage, even though the final thresholding process introduces some artifacts. A closer inspection reveals that the false positives in the thresholded image follow the path of the curvilinear structures in the original and enhanced image. The different filters and kernel dimension parameters used in this work had their unique side-effects in emphasizing certain curvilinear structures; their combination through scalar multiplication was intended to reduce the over-emphasis of those structures by individual filters. While the filters used significantly reduced the effect of Curvilinear structures in Figure 6 as well as the other images, they are nevertheless still pronounced in this image, which could be the cause of the false positives. The challenges encountered during the determination of an optimal threshold include mammogram image contrast, breast density and curvilinear structures. These are common challenges that

TABLE II COMPARISON WITH RELATED WORKS USING THE MIAS DATABASE

Author	Sensitivity (TP rate)	Specificity
Oh et al. [14]	93.1	87.5
Jian et al. [10]	83	
Mohanlin et al. [13]	96.55	60
Our work	100	11

are not unique to this project [14]. The results verify that intensity alone is insufficient as a criteria for the proper segmentation of microcalcification objects.

Table II shows a comparison with related works. The proposed model performs significantly better than related works in terms of sensitivity rates with regards to Normal/Abnormal ROIs with the highest score of 100%. However, it does not compare well in specificity rates, scoring 11%.

Looking at the sensitivity value in Table II, it can be concluded that the integration of the wavelet filters and the Laplace operators definitely contributed to the high detection rate for microcalcifications where present. This maximized the probability of positive pixels being identified, even if at the cost of falsely flagging non-calcification objects as positive. In practice, false negatives are more highly penalized than false positives [12]; false negatives imply delayed treatment as the alert is not raised, which might lead to a fatal prognosis for the patient. This supports the merits for this model in the sense that, its high sensitivity performance implies that calcifications are highly unlikely to be missed - the prompting of suspicious regions can be useful in contexts where such information is needed.

The strength of this model lies in its ability to detect all calcification objects where present. Regardless of the low specificity rate, the detection performance of the proposed model is nonetheless significant and can prove useful if used in the detection phase of the pipeline for calcification-related applications; future works can consider techniques such as incorporating user feedback rounds and refined tracing and elimination of curvilinear structures to possibly further refine the model in terms of specificity performance. The latter would have a major impact since most false positives were located along curvilinear structures. However, the scope of this work was to maximize the detection of microcalcifications as a preprocessing activity, which was sufficiently achieved.

#### IV. CONCLUSION

This study focused on fully automated detection of microcalcifications in digital mammogram images. It also investigated the discrimination between Malignant and benign subclasses of microcalcifications. To this end, the wavelet and Laplace filters were optimally integrated to amplify microcalcifications followed by postprocessing to reduce the number of false positives. Each of the filters have their unique strengths and side-effects in detection of calcificationlike objects; the essence was to combine them optimally to highlight their strengths while canceling their side-effects. The combined filter model detected all present calcifications, with a sensitivity rate of 100%, in all mammograms as demarcated by expert radiologists based on accompanying ground truth. The false-positive rate was significantly higher based on the lower specificity rate, a factor that can be investigated through use of efficient feature extraction and classification methods to characterize calcifications and curvilinear structures to reduce the false-positive rate. In accordance with the objectives of the study, this work has demonstrated the effectiveness of combining the likelihood maps from different filters in improving detection of calcification objects.

#### REFERENCES

- [1] V. Alarcon-Aquino, O. Starostenko, J. M. Ramirez-Cortes, R. Rosas-Romero, J. Rodriguez-Asomoza, O. J. Paz-Luna, and K. Vazquez-Muoz. Detection of microcalcifications in digital mammograms using the dualtree complex wavelet transform. *Eng. Int. Syst.*, 17(1):49–63, 2009.
- [2] Jacob Beutel and Milan Sonka. Handbook of Medical Imaging: Medical image processing and analysis, volume 2. SPIE Press, 2000.
- [3] Joseph D. Bronzino. *The Biomedical Engineering Handbook*, volume 2. CRC-Press, 2000.
- [4] Bhavana Deshpande, H.K. Verma, and Prachi Deshpande. Fuzzy based median filtering for removal of salt-and-pepper noise. *International Journal of Soft Computing and Engineering*, 2(3):76–80, 2012.
- [5] Y. G. Garud and N. G. Shahare. Detection of microcalcifications in digital mammogram using wavelet analysis. *American Journal of Engineering Research*, 02(11):80–85, 2013.
- [6] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison-Wesley, 2002.
- [7] N. B. Hamad, A. Masmoudi, and K. Taouil. Wavelets study for better multiresolution analysis in cad of microcalcification. *International Journal of Computer Science*, 10(1):372–378, 2013.
- [8] N. B. Hamad, K. Taouil, and M. S. Bouhlel. Mammographic microcalcifications detection using discrete wavelet transform. *International Journal of Computer Applications*, 64(21):17–22, 2013.
- [9] F.K. N. Harirchi, P. Radparvar, H.A. Moghaddam, F. Dehghan, and M. Giti. Two-level algorithm for mcs detection in mammograms using diverse-adaboost-svm. In *Pattern Recognition (ICPR), 2010 20th International Conference on*, pages 269–272, 2010.
- [10] W. Jian, X. Sun, and S. Luo. Computer-aided diagnosis of breast microcalcifications based on dual-tree complex wavelet transform. *BioMedical Engineering OnLine*, 2012.
- [11] Pelin KUS and Irfan KARAGOZ. Detection of microcalcification clusters in digitized x-ray mammograms using unsharp masking and image statistics. *Turkish Journal of Electrical Engineering and Computer Sciences*, 21:2048–2061, 2013.
- [12] Simon Marcellin, Djamel-Abdelkader Zighed, and Gilbert Ritschard. Detection of breast cancer using an asymmetric entropy measure. In Alfredo Rizzi and Maurizio Vichi, editors, *Computational Statistics* (COMPSTAT 06), volume XXV of Computational Statistics, pages 975– 982, Heidelberg, Germany, 2006. Springer. On CD.
- [13] B. Mohanalin, P.K. Karla, and N. Kumar. A novel automatic microcalcification detection technique using tsallis entropy and a type ii fuzzy index. *Computers and Mathematics with Applications*, 60:2426–2432, 2010.
- [14] Whi-Vin Oh, KwangGi Kim, Young-Jae Kim, HanSung Kang, JungSil Ro, and WooKyung Moon. Detection of microcalcifications in digital mammograms using foveal method. J Kor Soc Med Informatics 2009, 15(1):165–172, 2009.
- [15] M. Rizzi, M. D'Aloia, and B. Castagnolo. Computer aided detection of microcalcifications in digital mammograms adopting a wavelet decomposition. *Integrated Computer-Aided Engineering*, 16(6):91–103, 2009.
- [16] S.Abinaya, R.Sivakumar, M.Karnan, Murali Shankar, and Dr.M.Karthikeyan. Detection of breast cancer in mammograms - a survey. *International Journal of Computer Application and Engineering Technology*, 3(2):172–178, 2014.
- [17] Deepa Sankar and Tessamma Thomas. A new fast fractal modeling approach for the detection of microcalcifications in mammograms. J. Digital Imaging, 23(5):538–546, 2010.

- [18] P.D. Sathya and R. Kayalvizhi. Optimum multilevel image thresholding based on tsallis entropy method with bacterial foraging algorithm. *IJCSI International Journal of Computer Science Issues*, 7(5):336–343, 2010.
- [19] J. Suckling, J. Parker, D. R. Dance, S. Astley, I. Hutt, C. R. M. Boggis I. Ricketts, E. Stamatakis, N. Cerneaz, S.-L. Kok, P Taylor, D. Betal, and J. Savage. The mammographic image analysis society digital mammogram database. In *Proceedings of the 2nd International Workshop on Digital Mammography*, pages 375–378, 1994.
  [20] T.Balakumaran, I.Vennila, and C.G. Shankar. Detection of microcalcifi-
- [20] T.Balakumaran, I.Vennila, and C.G. Shankar. Detection of microcalcification in digital mammograms using one dimensional wavelet transform. *International Journal of Computer Science and Information Security*, pages 99–104, 2010.
   [21] T.Balakumaran, I.Vennila, and C.G. Shankar. Detection of microcal-
- [21] T.Balakumaran, I.Vennila, and C.G. Shankar. Detection of microcalcification in mammograms using wavelet transform and fuzzy shell clustering. *International Journal of Computer Science and Information Security*, 7(1):121–125, 2010.
- [22] Chia-Hung Wei and Chang-Tsun Li. Calcification descriptor and relevance feedback learning algorithms for content-based mammogram retrieval. In *Digital Mammography Lecture Notes in Computer Science*, volume 4046, pages 307–314, 2006.

2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech) Port Elizabeth, South Africa, November 26-27, 2015

# An Intelligent Locator Inventory to Accommodate Mass Customization in an Automated Flexible Fixture System

Andrew Illidge

Master's Student: Mechanical Engineering University of KwaZulu-Natal Durban, South Africa andrew.illidge1991@gmail.com Anthony Walker Lecturer: Mechanical Engineering University of KwaZulu-Natal Durban, South Africa Walkeraj@ukzn.ac.za

Glen Bright Professor: Mechanical Engineering University of KwaZulu-Natal Durban, South Africa Brightg@ukzn.ac.za

Abstract — The demand for mass customization is increasing, and while Flexible Manufacturing Systems provide customization, they do not currently meet the requirements of mass customization. The main limitation is in the fixtures that are utilized. They generally do not provide rapid setups and setup changeovers, accommodate a variety of part families and geometries and incorporate Artificial Intelligence in the setup and fixture planning phases. There is therefore a need for a fixture system that meets these requirements if mass customization is to be met. This paper aims to present an Automated Flexible Fixture System that utilizes an intelligent locator inventory to meet the requirements of mass customization.

Keywords—Flexible Manufacturing Systems; Flexible Fixtures; Group Technology

#### I. INTRODUCTION

The demand for mass customization is growing and with the high costs involved in fixture design and manufacture, a fixture that is capable of handling a variety of different part families and geometries is required. This demand has partly motivated the development in Flexible Manufacturing Systems (FMS) [1]. Whilst FMS are useful, a major bottleneck hindering mass customization has been the fixtures that are utilized [2]. Along with the design and manufacture costs of fixtures, the time required to perform setup changeovers and to learn the setup procedures for new parts can be high. The use of a fixture that can automatically accommodate a variety of different part families and geometries, perform setup changeovers automatically and rapidly, and minimize the time required to learn new fixturing procedures is necessary if mass customization is to be facilitated. What was described is what is known as an Automated Flexible Fixture System (AFFS) [1].

Automatically accommodating a variety of part families and geometries requires the use of a number of locators in the form of intelligent locator inventory and a locator selection algorithm to ensure that the most appropriate locator is chosen. Along with a locator selection algorithm, a setup planning algorithm and fixture planning algorithm are required. This encompasses two of the four phases of fixture design, setup planning and fixture planning. This paper aims to present a design for an AFFS and the expert system algorithms that were used in locator selection and fixture design.

#### II. LITERATURE REVIEW

#### A. Fixtures

A fixture is device that is used to constrain a workpiece in order to resist forces induced by machining processes [3]. A fixture differs from a jig in the way that it presents the workpiece to the cutting tool. A fixture keeps the workpiece stationary relative to the moving cutting tool, whilst a jig moves the workpiece relative to the stationary cutting tool. The key requirements of a fixture are to allow for repeatability, minimize part deformation and reduce vibrations [3]. A typical fixture consists of a locator, clamps and a baseplate [4]. The locator resists the primary cutting forces and positions the workpiece so that repeatability can be achieved. The clamps resist the secondary cutting forces, while pushing the workpiece against the locator, and the baseplate is the platform upon which the locator and clamps are attached [5]

There are two categories of fixtures, being dedicated and flexible fixtures. Dedicated fixtures are designed and manufactured to constrain a single workpiece, while flexible fixtures are designed and manufactured to constrain multiple workpieces. [2] There are three types of flexible fixtures: modular fixture; reconfigurable fixtures; and conformable fixtures [2]. A modular fixture consists of a fixture part inventory. The parts can be combined to form numerous arrangements, with each arrangement being constructed to accommodate a specific workpiece. The disadvantages to using modular fixtures lye in the uncertain tolerance that is involved with combining fixture parts with different tolerances, the somewhat limited amount of fixture combinations that can be achieved and the storage space needed to accommodate the fixture part inventory [2].

A reconfigurable fixture can be adjusted to accommodate workpieces of variable geometries. The disadvantage of using reconfigurable fixtures is the fact they are generally limited in both the number of part families and the amount of workpiece geometries that they can constrain, more so then modular fixtures.

A conformable fixture, as it name suggests, conforms to the shape and geometry of the workpiece [2]. There are two main types, pin-matrix and phase-change material [2]. These fixtures provide the most flexibility, as they can conform to the surface of most part geometries and families [2]. The issue facing conformable fixtures is in their use as locators and their cost. Their use as locators is difficult to implement, as it requires a reasonable amount of locator intelligence to determine the workpieces position based on the shape of the conformed fixture.

#### B. Fixture Design

The fixture design process comprises four phases; setup planning, fixture planning, unit design and verification. The setup planning phase involves grouping machining operations that can be done in a single setup. The aim is to minimize the number of setups used, as this reduces the Manufacturing Lead Time (MLT) [3]. The fixture planning phase addresses six fixture requirements: constraining (stability and deformation); physical (type of locators and clamps to be used); tolerance; collision prevention; usability; affordability [3]. It also involves the fixture analysis and synthesis processes. Fixture analysis is the process of modelling the fixture planning requirements to the available fixture variables [1]. Fixture synthesis is the process of using the models developed in the fixture analysis process to select variable that meet fixture planning requirements [1]. The unit design phase involves both conceptual and physical design and aims to develop locators and clamps that meet the fixture planning requirements and configuration [3]. The verification phase aims to verify that the fixture parts designed in the unit design phase meet the requirements of the fixture planning phase [3].

#### C. Computer-Aided Fixture Design (CAFD)

CAFD involves using computer software to perform the fixture design process. The software that is used generally focuses on a single fixture design phase, rather than integrating all four phases [2]. This, and the need for human interaction, are the major issues that CAFD faces [2][6]. The phases into which most of the research has been conducted are the fixture planning phase and verification phase [2].

In fixture planning, there are numerous methods used in the fixture analysis and synthesis processes in order to create a fixture layout (locator and clamp configuration). In fixture analysis the following methods are used to model the fixture requirements to the workpiece: a geometry approach, which uses geometric data from the CAD model to determine possible clamping and locating surfaces; screw theory, which is a mathematical approach that determines the minimal amount of clamps and locators required to constrain the workpiece; and a rule-based approach that utilizes feature data from the CAD model to determine possible locating and clamping positions [1][2]. In fixture synthesis the following methods are used: an expert system, which is a heuristic reasoning approach based on a knowledge database; Case-Based-Reasoning (CBR) which utilizes previous workpieces', of a similar part family and geometry, fixture layouts as a blueprint for the current fixture layout; and Genetic Algorithms (GA) to achieve an optimized fixture layout [1][2].

#### D. Automated Flexible Fixture Systems

An AFFS is a fixture system that ideally can automatically perform the fixture design process, automatically perform setup changeovers, automatically initiate clamping and locating configurations, accommodate a variety of different part families and geometries, provide adaptive behavior and have a form of Artificial Intelligence [1].

The use of flexible fixtures in FMS has aided in cost reduction by reducing the required amount of dedicated fixtures that must be designed and manufactured for each new part [2]. Whilst this is useful for job shop production, their capacity for mass customization, individually, is limited. A hybrid flexible fixture, that utilizes modular, reconfigurable and conformable fixtures, is required if mass customization is to be achieved in AFFS.

#### E. Group Technology (GT)

Locating a workpiece requires low tolerances in order to ensure that repeatability is achieved. Locators must therefore be well suited to the workpiece that they are locating. Just as GT has improved the manufacturing environment in areas such as automated process planning, materials handling, production and inventory control and product design, it may be used to group and map part features to specific locators in order to simplify the fixture design process [7][8]. The use of Group Technology in the form of a locator inventory could therefore aid the accommodation of a variety of part families and geometries.

#### III. AUTOMATED FLEXIBLE FIXTURE SYSTEM DESIGN

An AFFS was designed that meets the following requirements: automated setup changeover; accommodation of a variety of part families and geometries; semi-automatic CAFD; adaptive behavior in the form of clamping force control; and automatic locating and clamping. It consists of the following components: a locator inventory; a two DoF (rotational) serial manipulator for workpiece positioning; two electronic pin matrix clamps; two four DoF (three linear and one rotational) clamp actuators; and a T-slot baseplate. The CAD design assembly can be seen in Fig. 1.



Figure 1: AFFS CAD Model Assembly

#### A. Automatic Setup Changeovers and Constraining

The use of a two DoF workpiece positioning mechanism and two four DoF clamp actuators enables automatic setup changeovers and automatic clamping and locating. Positioning data is sent to stepper motors which provide the actuation of both the workpiece positioning mechanism and clamp actuators through either lead screw drives (where linear translation is required) or gearboxes (where rotational movement is required). The process that the AFFS follows in order to constrain a workpiece can be seen Fig. 2.



Figure 2: AFFS Process Flow Diagram

### *B.* Accommodation of a Variety of Part Families and Geometries

Pin matrix clamps conform to the shape of the surface to which they are pressed, and can then lock the conformed shape in place. This allows for the clamping of a variety of part families and geometries. The pin field also reduces the stress induced in the pins and workpiece due to its force distribution, and so the unit design phase of fixture design can somewhat be neglected. The pin matrix CAD model design can be seen in Fig. 3.



Figure 3: Pin Matrix Clamp CAD Model

An intelligent locator inventory is used whereby locators are mapped to part features. The locators can be rapidly interchanged based on workpiece requirements and form the modular part of the AFFS. Each locator, according to its purpose, is intelligent in some way. For instance, in the case of the two-finger parallel gripper locator, the finger gap size can be adjusted (using a servo motor) and recorded to accommodate a variety internal hole gap sizes. The locator inventory can be seen in Table 1.

Table 1: Locator Inventory

	Locator Inventory				
1	3-Jaw Chuck locator				
2	Pin field locator				
3	2-Finger angular gripper				
4	2-Finger parallel gripper				
5	Vise locator				
6	Cylindrical pins				
7	Feature pins				

#### C. Adaptive Behaviour

Each pin matrix clamp utilizes a force sensor to determine the clamping force that is being exerted. A vibration sensor inside each locator provides feedback and adaptivity is displayed in its ability to reduce vibration by adjusting the clamping forces.

#### D. Semi-Automatic CAF

Using feature recognition software, the CAD model can be broken down into features. For instance, a cylindrical cavity is classified as a circular hole. A geometric analysis determines where the feature is relative to a reference point on the model. Features are grouped based on the type of machining process they require and on their position relative to one another. Both feature recognition and geometric analysis are done using Solidworks CAD software.

Using the data obtained from the feature and geometric analyses, an expert system is used to select a locator, that best suites the workpiece, and to perform the setup planning, fixture analysis and synthesis processes. Three algorithms were constructed for the expert system, a locator selection algorithm, a setup planning algorithm and a fixture analysis and synthesis algorithm. The locator selection algorithm, as seen in Fig. 4, uses the geometric and feature data extracted from the CAD model of the workpiece to select a locator from the locator inventory to be used in the AFFS. The locator inventory has a hierarchy and therefore in the case that multiple locators fit the workpiece, the locator with the highest ranking is selected. The hierarchy order is based on how well the locator collaborates with the AFFS, as well as its location simplicity (the ease with which it locates the workpiece). There however criteria under which a locator with a lower ranking maybe selected over one with a higher ranking, when either one may have been used. For instance, if the internal features lie within 15mm over each other, a pin field locator or vise locator is selected over the 2-finger parallel gripper locator.

The setup planning algorithm, as seen in Fig. 5, maps features to one of six setups. The six setups are based on the type of locator that is used and do not contain the exact clamping coordinates but rather an approximated clamping configuration. These approximated setups are constant to the locator and do not differ from workpiece to workpiece. Features that occupy the same (or similar) plane or surface are grouped together to form a setup.







Figure 5: Setup Planning Algorithm

The fixture analysis and synthesis algorithm, as seen in Fig. 6, selects the exact clamping coordinates for each of the six setups. Surfaces or areas of the workpiece where minimal machining takes place and that contained sharp angles and already machined distinct features were targeted as clamping options. The use of distinct features as clamping positions creates rigidity and was possible due to the use of pin matrix clamps, as they can conform to most surfaces. Along with areas of minimal machining, areas that contained the entirety of a feature group are targeted. This was done in order to allow setup changeover and tool change to coincide where possible.



#### Figure 6: Fixture Planning Algorithm Flow Diagram

#### IV. CONCLUSION

Mass customization poses a variety of complications, with a key one being the fixtures that are utilized. Generally FMS cannot facilitate mass customization due to the inflexibility of their fixtures. In the case that flexible fixtures are used, the time required to learn and perform setups and setup changeovers for a variety of different part families and geometries can be substantial and therefore limits the viability of mass customization.

It is therefore necessary to implement an AFFS in FMS if mass customization is to be facilitated. The AFFS must meet the following requirements: accommodate a variety of part families and geometries; automatically perform setups and setup changeovers; provide adaptive behavior to reduce vibration during machining; and offer intelligence to automatically (or semi-automatically) generate setups and clamping positions. The AFFS that was designed meets these requirements to a certain degree, however further work in manufacture, software design and testing is required before it is fully operational. A more detailed set of algorithms and coding is required before the semi-automated CAFD is useful.

#### REFERENCES

- [1] Z. M. Bi and W. J. Zhang, "Flexible Fixture Design and Automation: Review, Issues and Future Direction," *Int. J. Prod. Res.*, vol. 39, no. 13, pp. 2867–2894, 2001.
- [2] H. Wang, Y. Rong, H. Li, and P. Shaun, "Computer Aided Fixture Design: Recent Research and Trends," *Comput. Des.*, vol. 42, pp. 1085–1094, 2010.
- [3] I. Boyle, Y. Rong, and D. C. Brown, "A Review and Analysis of Current Computer-aided Fixture Design Approaches," *Robot. Comput. Integr. Manuf.*, vol. 27, pp. 1–12, 2011.
- [4] Y. Zheng and W.-H. Qian, "A 3-D Modular Fixture with Enhanced Localization Accuracy and Immobilization Capability," Int. J. Mach. Tools Manuf., vol. 48, pp. 677–687, 2008.
- [5] E. Hoffman, *Jig and Fixture Design*, 5th ed. Delmar, Cengage Learning, 2004, pp. 8–48 of 359.
- [6] X.-T. Yan and B. Eynard, Global Design to Gain a Competitive Edge: An Holistic and Cllaborative Design Approach based on Computational Tools, 1st ed. Springer-Verlag London Limited, 2008, pp. 765–767.
- [7] H. S. Bawa, *Manufacturing Processes II*, 1st ed. New Delhi: Tata McGraw-Hill, 2004, pp. 265–268.
- [8] R. K. Rajput, A Textbook of Manufacturing Technology: Manufacturing Processes. Laxmi Publications (P) LTD, 2007, pp. 848–855.

# Using a Classifier to Track the Edge of a Conveyor Belt

Ra'eesah Mangera, Dr Garry Morrison, Anthon Voigt R&D De Beers Technologies SA De Beers Group Services Johannesburg, South Africa raeesah.mangera@debeersgroup.com

Abstract-Ensuring that a conveyor belt runs true on an Xray transmission (XRT) particle sorting system is important for a number of reasons. It limits the damage to the edge of the belt, reduces wear to the conveyor and minimises material spillage. Knowing where the belt is also defines the region of interest (ROI) for the applied algorithm, thus reducing false positives outside the ROI. In addition, correct belt tracking ensures that particles are all contained within the detectable area, avoiding the loss of material. Therefore, ensuring correct conveyor belt tracking can greatly impact the cost and efficiency of the sorting system. This paper presents a novel method for tracking the edge of a conveyor belt in a dual energy X-ray transmission (DE-XRT) sorting system. Using a classifier to determine where the belt edge is, the position of the belt is estimated with a pixel error of less than 15 pixels across all test images. An error of this order is expected as the belt edge is 18 pixels wide in the image. The current method, which is not affected by particles lying against the belt, is shown to be better that the previous approach using morphology, which exhibited poorer tracking accuracy when particles touch the belt. In addition, no additional sensors need to be installed and the belt does not need to be specially manufactured, reducing the cost of the conveyor system. Keywords-belt edge detection; Gaussian classifier; X-ray

#### I. INTRODUCTION

transmission; dual energy

In our dual-energy X-ray transmission (DE-XRT) sorting system material is fed onto a conveyor belt. The material is imaged using an X-ray source and the images are analysed in order to sort the material. Knowing where the belt is in the image defines a region of interest (ROI) for the sorting algorithm. In addition, correct belt tracking ensures that the detectable area contains all particles avoiding loss of material. Correct belt tracking also reduces wear to the belt and the conveyor system and minimises material spillage [1]. This paper presents a novel approach which uses a classifier to detect the edges of the belt in each frame, thus enabling belt tracking. Section I-A presents a brief overview of the operational principles of a DE-XRT system, Section II provides a review of other belt tracking methods, Section III presents the classifier approach employed. The results are discussed in Section IV and Section V concludes.

#### A. Dual Energy X-Ray Transmission (DE-XRT) System

A dual-energy X-ray transmission (DE-XRT) system makes use of a dual-energy X-ray line scan sensor to generate images of the transmitted X-rays, similar to those seen at



Fig. 1: Illustration of a dual-energy X-ray transmission (DE-XRT) system.

airport baggage security. DE-XRT systems allow for rapid approximation of the atomic number which is used to evaluate the mineral content of materials on a conveyor belt. Dualenergy refers to the camera that contains two sensors, one responding to low energy X-rays and one responding to high energy X-rays. Each material characteristically absorbs X-rays based on it's atomic structure and density and this absorption is reflected in the high and low energy camera signals. As there is a variation in the exposure response of the X-ray detectors, flat-field and dark-field corrections must be made. Offset correction is used to account for dark current effects and gain correction is used to correct for the non-uniform response of individual pixels. The conventional method to correct a raw X-ray image is given by Kwan, Siebert and Boone [2],

$$C(x,y) = \bar{g} \frac{I(x,y) - B(x,y)}{G(x,y) - B(x,y)}$$
(1)

where C is the corrected image (also known as the absorption image), I is the raw image, B is the averaged dark-field image, G is the averaged flat-field image and  $\bar{g}$  is the mean pixel value of G. Therefore using a DE-XRT system, two X-ray absorption images are obtained. The system is illustrated in Fig. 1.

#### **II. PREVIOUS WORK**

There are many commercial options available for tracking a conveyor belt, however the majority of these rely on the



Fig. 2: An example of a low energy profile for an X-ray transmission system frame.

installation of additional sensors or modification of the conveyor belt. For example the system from Tru-Trac [3] employs an internal pivot perpendicular to the plane of the belt with tapered ends that actuates a trainer when the belt moves off centre. Intralox use 90-degree sensor rollers at some point in the assembly to detect if the belt is off-centre [4]. The Phoenoguard [5] system from Phoenix requires the system to be installed on the belt. In the work presented by Balbin and Karmakar [6], RFID tags are installed on the conveyor belt and used for conveyor belt tracking. Zhu et al. [7] use a video camera and edge detection to track the edge of the belt. This is similar to the approach employed here, but instead of RGB images, X-ray images are used. Zhang and Lou [8] use a mechanical displacement sensor to calculate the belt deviation. All of the above approaches would require the installation of additional sensors whereas the classifier-based approach makes use of sensors already installed (i.e. the X-ray camera).

#### A. Previous approach

A morphological approach to belt tracking was previously employed. First, the raw low energy profile of the image is extracted. Fig. 2 is the mean of each column in the image. The periodic dip in signal is due to the fact that the sensor is made up of a series of camera boards. The dips are due to the edge-pixel effect, where edge pixels experience a lower signal. The number of pixels between each camera board is known, therefore we can eliminate the minima between camera boards by averaging the two pixels on either side of adjacent camera boards. Then, a filter is used to look for the shape of the belt edge in the profile (shown in Fig. 3). However, this approach is not robust in cases where there are particles touching the outside of the belt, as this resembles the belt edge profile.

#### III. DETECTING THE BELT EDGE

In order to set a region of interest for the classification algorithms we must know the location of the belt edge. The belt does not run true over time and will wander from



Fig. 3: Morphological operator used to detect the edge of the belt.



Fig. 4: Extreme example of material spillage. If the edge of the belt is not detected the spillage is classified in each frame. This may affect the results.

left to right. Additionally as the belt is not manufactured perfectly straight, during one rotation the edge will oscillate. For these reasons the region of interest can wander and must be detected on a frame-by-frame basis. Ensuring the region of interest tracks the belt ensures that only material on the belt is classified, i.e. material spillage is not classified. An extreme case of material spillage can be seen in Fig. 4. The objective of the study is to detect the ROI that robustly excludes areas outside of the belt.

#### A. Problem Description

Unlike other approaches that use an RGB image [7], the X-ray images do not have a clear line where the belt edge is located, therefore simple edge detection cannot be employed. However, we note that due to a shoulder at the ends of the belt, the belt is thicker at it's edge. This is illustrated in Fig. 5. As the X-ray absorption in the dual-energy space is related to the thickness of material we can use this to design a classifier to highlight pixels at the belt edge.

There are three cases in which the belt can operate, i.e. the edge may run left of the flat-field, right of the flat-field or exactly on the flat-field. In the first two cases, there would be a non-zero value at the belt edge (cases are shown in Fig. 6). The deviation from the flat-field can be used to train a classifier to identify the belt edges. In the third case, it appears that no belt edge would be detectable. However the flat-field is the mean of a few images, covering a full revolution of the belt. Hence the flat-field includes the side-to-side wobble of the belt. Therefore it is likely that these pixels will also record a detectable signal.







Fig. 6: Cases in which the belt can operate. (a) Belt is left of the flat-field. The pixels highlighted in purple will be classified as edge. (b) Belt is right of the flat-field. The pixels highlighted in purple will be classified as edge. (c) Belt is exactly aligned with the flat-field. All edge pixels will be classified as belt.

#### B. Classifier Design

As the belt edge cannot easily be seen in X-ray images with the naked eye, thumb tacks were placed on the left hand belt edge. These images were then used to manually paint the belt edge and belt classes e.g. Fig. 7. A scatter plot of an empty belt, with the edge pixels marked is shown in Fig. 8.

In total 21 images were captured. These images were captured over a period of 10 months whilst the machine was operating in the field. The images emulate three different scenarios, an empty belt, a belt with material on and lastly cases where there are particles touching the outside of the belt. Raw images showing these three cases can be seen in Fig. 9. A line scan camera is used, therefore when there are particles outside of the belt they appear as streaks through the image as the particles are stationary with respect to the belt.

Data from two of the images was used to train a Gaussian classifier for the edge pixels and another Gaussian classifier for the belt pixels. The two Gaussian classifiers are then cascaded



Fig. 7: Image where the different classes have been painted using the Perclass [9] toolbox. Magenta is the thumb tacks placed in the belt edge, the red is where the belt edge is located, green is belt and grey is the side of the belt. The blue represents an 'unknown' class.



Fig. 8: Absorption values of belt pixels and belt edge pixels. Thumb tacks place in the belt edge is used to localise the belt edge. Then the edge and belt pixels are manually marked.

to determine whether a given pixel is an edge, belt or particle. The designed classifier can be seen in Fig. 10. A sample of the classification on a test image can be seen in Fig. 11. Note that some of the particle pixels are classified as edge. However since we expect the belt edge to be at the sides of the image, only the first and last 100 pixels are classified.

#### C. Edge Detection

As can be seen in Fig. 11 the edge is detected along multiple columns (indicated in red). This is expected as the belt edge is 18 pixels wide. Therefore to obtain a single pixel location for the belt edge on either side, we consider the belt edge to be located in the column with the most amount of pixels classified as edge. The overall system design is shown in Fig. 12.



Fig. 11: Example of the classification on a test image. Green represents pixels classified as belt, red pixels classified as belt edge and blue pixels classified as particle.



Fig. 12: Overview of the belt edge detection system.

#### IV. RESULTS

Using the remaining 19 images with thumb tacks, the edge of the belt was manually estimated. An example of this is given in Fig. 13. This is considered to be the ground truth for each image. Comparing the classifier result to the manual ground truth result enabled the pixel error to be calculated for each edge of each image. The difference between the true belt edge and the detected belt edge is less than 15 pixels in all cases, as seen in Table I. This is considered to be acceptable as the belt edge actually covers 18 pixels. Therefore we expect an error between the actual and detected edge of this order. On average, the pixel error is approximately 7 pixels. Even in images where there are particles on the outside of the belt, such as image number 16 and 18, the error is less than 4 pixels on the left side where there is a particle. Using the previous approach this error was 29 pixels. Therefore the classifier method is robust to cases where there are particles on the outside of the belt. In this scenario the morphology-based method incorrectly identified the belt edge due to the particle profile.

The confusion matrix for the classifier can be seen in Table II. Note that all belt pixels are correctly classified but only 56% of the edge pixels are classified as edge. This is due to the large overlap between the two classes as can be seen in Fig. 14. Note that in the high energy there is a large proportion of the edge class and belt class are inseparable, therefore we expect there to be confusion between the classes. This can be explained by the amount of overlap between the flat-field and

TABLE I: Pixel error on the left and right sides of the belt for each image in the dataset

Image Number	Left Error (pixels)	Right Error (pixels)
1	9	2
2	15	1
3	11	2
4	10	1
5	10	1
6	10	1
7	13	0
8	15	2
9	13	1
10	10	0
11	13	0
12	10	0
13	1	3
14	9	15
15	1	14
16	5	12
17	5	8
18	3	6
19	1	5
20	10	9
21	6	13

current image, as illustrated in Fig. 6. Only pixels which do not overlap the flat-field are classified as edge. However, for this application as long as there is at least a single column offset, the belt edge will be detected.



(a) Raw high energy image of an empty belt.



(b) Raw high energy image for the case where there is material on the belt.



(c) Raw high energy image for the case where there is spillage on the left-hand side of the image, highlighted by the red block.

Fig. 9: Examples of images captured.

#### V. CONCLUSIONS

This paper presented a novel method for tracking the edge of a belt using a classifier in a DE-XRT system. The method was shown to be robust for a wide range of cases from an empty belt, to cases where there is spillage. An advantage over other belt tracking systems is that no additional sensors need to be installed to detect the edges of the belt.



Fig. 10: Classifier used to detect the edge of the belt in dualenergy feature space. The red dots are edge pixels from the test set, the green dots are belt pixels from the test test. The different coloured regions in the background represent the classifier decisions, with blue being edge, yellow belt and red particle. Note that the red points in the red coloured region will be misclassified as particle.



Fig. 13: An example of a manually marked ground truth image using thumb tacks as a guide.

#### A. Future Work

As mentioned previously, the edge is localised with an average pixel error of approximately 7 pixels. Further work could be done to reduce this error, for example instead of using the column with the most number of edge pixels, the left-most or right-most column could be used. In addition, other classifiers could be investigated to detect the belt edge.

TABLE II: Confusion matrix for the belt edge classifier.

		Decisions		
		Belt	Edge	Particle
	Belt	100	0	0
True Labels	Edge	44	56	0
	Particle	0	11	89



Fig. 14: Distribution of the training data for the belt and belt edge classes in the high energy and low energy spaces.

#### REFERENCES

- L. Zhao and Y. Lin, "Typical failure analysis and processing of belt conveyor," *Procedia Engineering*, vol. 26, pp. 942–946, 2011. [Online]. Available: http://dx.doi.org/10.1016/j.proeng.2011.11.2260
- [2] A. L. C. Kwan, J. A. Seibert, and J. M. Boone, "An improved method for flat-field correction of flat panel x-ray detector." *Medical physics*, vol. 33, no. 2, pp. 391–393, 2006.
- [3] Tru-Trac, "Tru-Trac Trough Tracker conveyor belt tracking system," 2007. [Online]. Available: http://www.tru-trac.com/trough\_tracker.htm
- [4] Intralox, "Activated Roller Belt 90-Degree Transfers," 20111. [Online]. Available: http://www.intralox.com/90-degree-transfers.aspx
- [5] Phoenix Conveyor Belts, "Phoenoguard PX," 2013. [Online]. Available: http://www.phoenix-conveyorbelts.com/pages/products/protectionsystems/phoenoguard-px/phoenoguard-px\_en.html
- [6] I. Balbin and N. Karmakar, "Novel chipless RFID tag for conveyor belt tracking using multi-resonant dipole antenna," European Microwave Week 2009, EuMW 2009: Science, Progress and Quality at Radiofrequencies, Conference Proceedings - 39th European Microwave Conference, EuMC 2009, no. October, pp. 1109–1112, 2009.
- [7] A. C. Zhu, W. Hua, C. Wang, and Y. X. Wang, "Research on the measurement of belt speed by video in coal mine based on improved template matching algorithm," *Journal of Coal Science and Engineering*, vol. 17, no. 4, pp. 469–474, 2011.
- [8] J. W. Zhang and P. H. Lou, "Automatic Detection and Hydraulic Correction Technology of Belt Deviation," *Mechanical Engineer*, vol. 10, p. 13, 2008.
- [9] P. Paclik and C. Lai, "perClass: Practical Machine Learning," 2006. [Online]. Available: http://www.perclass.com

## Developing a Java based RFID application to automate student attendance monitoring

R.B. Kuriakose

Department of Electrical, Electronics and Computer Engineering, Central University of Technology, Private Bag X20539,Bloemfontein, South Africa, 9300. Email: rkuriako@cut.ac.za

*Abstract*- Research has proved that there is a high correlation between high class attendance and academic performance. Learning is a progressive activity which builds upon those of the previous day(s). Reading material and working independently do not compensate for the loss of insight gained by being physically present in classroom environment. As a result attendance monitoring should be an integral part of any institution especially a tertiary institution.

The major challenge that a facilitator has in monitoring attendance is the sheer paper work and subsequent administrative work that follows the process if it is done manually. Many techniques have been developed to automate attendance monitoring including the use of biometrics. These techniques have their own strengths and weaknesses.

This paper describes a pilot experiment that was rolled out at the Central University Of Technology, Free State, South Africa where a Java based Radio Frequency Identification Technology (RFID) application was used to automate attendance monitoring.

The paper first compares the challenges facing manual attendance monitoring. Secondly it looks at some of the possible solutions for automating attendance monitoring including RFID technology. Thirdly it shows how the hardware part is setup and finally it describes how Java programming is used to create a meaningful link between the data collected by the RFID hardware so as to automate the attendance register.

*Keywords—RFID, attendance monitoring, Java programming, tags, RFID reader* 

#### I.INTRODUCTION

Truancy is a serious problem that affects almost all spectrums of higher education in numerous universities[1] across the world. Study also states that there is a direct correlation[2]–[4] between class attendance and academic performance or non-academic performance. A study done at the Central University of Technology also shows that poor academic performance may lead[5] to students dropping out of the university altogether.

Mandatory attendance monitoring for each and every class is seen as vital intervention to combat truancy [1]. At the H.J. Vermaak

Department of Electrical, Electronics and Computer Engineering, Central University of Technology, Private Bag X20539, Bloemfontein, South Africa, 9300. Email: <u>hvermaak@cut.ac.za</u>

Central University of Technology, currently, attendance for each class is taken by manually signing on to an attendance register. This paper firstly looks at the disadvantages of manual attendance monitoring, secondly it analyses the different techniques currently used in automating attendance monitoring and finally proposes how Radio Frequency Identification (RFID) technology in conjunction with Java programing is used as a possible alternative to automate the attendance register.

### II.CURRENT SYSTEM FOR REGISTERING ATTENDANCE AND ITS DISADVANTAGES

As described in the introduction, currently attendance monitoring is done manually. In this scenario, the facilitator has a list containing the names as well as student numbers of all students registered to do a specific course. The list is circulated during a class session among the students. Students are supposed to sign next to their name and student number. This process is repeated for each session for the entire semester. At the end of the semester the facilitator has to manually transfer the data on to an excel sheet and tabulate the attendance in percentage for each student.

This section explains some of the many challenges that a facilitator experiences as a result of manual attendance monitoring. They show how this often results in the complete process of attendance monitoring being undermined especially if the class strength is very high.

- Time consuming and tedious
- Loss of data and neglect
- Error prone
- Difficult to monitor

#### Time consuming and tedious

The primary responsibility of a facilitator in higher education[6] is to provide scientific knowledge to students such that they are able to produce solutions for the problems facing the society and country as a whole. Attendance monitoring becomes a time consuming and tedious procedure when it is done manually. This starts right from drawing up the list to the time the attendance is calculated on the excel sheet. It needs to be stressed that this work is often on top of all other administrative work that a facilitator has to complete in a semester.

#### Loss of data and neglect

The sheets of paper collected for the purpose of monitoring attendance manually can easily be misplaced or lost from the facilitator before they are recorded on to a computer. It might sometimes also be the case that the facilitator forgets to take the attendance list to class, this might result in neglecting the process altogether.

#### Error prone

The attendance list circulating though the class might be a hindrance to students who are otherwise focussed on the lecture. This might result in the list passing without their signatures or signing on a row which is not assigned to them. This scenario is prevalent in most first year classes at the university where the numbers are quite high.

#### Difficult to monitor

The main intention of attendance monitoring is to encourage class attendance by all the students. In an average eighty minute class with a class size over one hundred students it is practically impossible to ensure if the right student has entered the name in the right column or if the students have acted as proxies for their fellow classmates.

### III.TECHNOLOGIES AVAILABLE FOR AUTOMATING ATTENDANCE REGISTER

The disadvantages of manual attendance monitoring have been documented in Section II. The obvious solution to overcome the challenges raised by the current system is to automate the student attendance register. This section compares three existing technologies that may be used for this purpose. They are;

- Barcode systems
- Biometric systems
- RFID systems

The comparison shown in Table 1 is made by considering parameters such as data density, cost, data reading speed and influence of interference between reader and data carrier. These parameters are considered important if automating the student attendance register.

Barcode systems are the most common automating technology available. They are considerably cheaper than Biometric systems and RFID systems. A barcode system consists of a data carrier and a data reader. The data carrier is a unique code which can be printed and pasted on to a device that needs to be scanned. The data reader is an optical scanner which picks this data and a computer can interpret this data such that it makes sense to the user. The major disadvantage of barcode systems is the influence of dirt or covering the data carrier. This is especially critical in a student environment.

The Biometric systems would have been the ideal solution as it is widely used in attendance registration currently[7]. The most common biometric scanner is the fingerprint scanner. The fingerprints of the users are initially stored in the scanning device. Once the scanner is used, the fingerprint of the user is scanned with the stored data and if a match is found it will be recorded with a time stamp, there by registering the user. The purchasing and operating cost of these devices is currently quite high. If each lecturer in a university environment has to be given a Biometric system for attendance registration, this would make the project very expensive. This is its major drawback.

The last of the technologies compared for this study, looks at RFID technology. It can be seen from the comparison in Table 1 that RFID technology is well suited for automating the student register as compared to the other technologies that are available with respect to the parameters. Section IV looks extensively at this technology.

TABLE I: COMPARISON BETWEEN DIFFERENT AUTOMATIC IDENTIFICATION SYSTEMS

Parameters	Barcode system	Biometric systems	RFID systems
Data density	Low data density	High data density	Very high data density
Influence of covering the data carrier	Total failure of system	Total failure as system works on contact	No influence
Influence of direction between reader and data carrier	Failure - if no line- of-sight communi cation	Not applicable as direct contact is needed	No influence as data are transferred via radio waves
Purchasing and operating cost	Low	High	Low
Distance between reader and data carrier (in centimeters)	0-50 cm	Direct contact	0-6 m depending on the frequencies used

#### **IV.RFID TECHNOLOGY**

This section elaborates on RFID technology. An RFID system consists of three main components, RFID tags, RFID readers and RFID middleware.

#### RFID tags

The RFID tags or transponders consist of a chip and an antenna[8]. A chip or tag can store a unique identification number (UIN) or some other form of data which differentiates it from another tag. The data storage capacity of a RFID tag,

cost, transmission range and type of memory depends on the type of RFID tag that is used. There are three types[9] of RFID tags. They are passive, semi passive and active tags. A comparison between the different tags based on the parameter mentioned in this section is shown in Table 2.

TABLE II: COMPARISON BETWEEN DIFFERENT TYPES OF RFID TAGS

Parameters	Passive tags	Semi Passive	Active tags	
		tags		
Transmission	Short (up to 6m)	Medium (up	Long (up to	
range		to 30m)	200m)	
On board power	No (powered by the	Yes (Internal	Yes (Internal	
supply	reader)	battery	battery	
Memory	Mostly read only	Read and	Read and	
		write	write	
Cost	Cheap	Medium	Expensive	

The transmission range of an RFID tag is depended on the types of tags as shown in Table 2. The three tags operate [8] on three frequencies; they are low frequency, High frequency and Ultra high frequency. The low frequency (LF) passive RFID tags operate at 125-134 kHz and have a range of up to 6cm. The high frequency (HF) Semi-passive RFID tags operate at 13.56MHz and also have a range of up to 20cm. The Ultra high frequency (UHF) Active RFID tags have a range of above 3m up to 200 m and operate at 868-928MHz.

The main difference between the LF and HF frequencies is that the HF RFID tags are often packaged in foil inlay and have a higher data read rate than low frequency tags.

#### **RFID** readers

The RFID readers are scanning devices[10]that are used to retrieve information from the RFID tags. A typical RFID reader has two functional blocks[8]. They are; the high frequency (HF) interface and the Control unit.

The HF interface[11] is the main part of the RFID reader. It is responsible for powering the passive tag, modulating the signal send to the RFID tag and receiving the data from the RFID tag.

The control unit of the RFID reader is the secondary part of the RFID reader, which is responsible[11] for communication and execution of the applications software commands, signal coding and decoding and communication control with a transponder.

#### RFID middleware

RFID middleware is the backbone of an RFID system [12]. It transforms RFID from a form of Automatic Identification and Data Capture (AIDC) technique to an intelligent technology by making sense of the data captured from the RFID tag[8]. It is essentially a software program written to work as a background thread to provide processed data to back end applications[13].

A RFID middleware has three basic layers that enable it to act as the interface between the RFID reader and the RFID applications. These are;

- Reader adapter
- Event manager

- Application level interface
- Middleware management

#### Reader adapter

This is the lowest layer[14] of the RFID middleware. The purpose of this layer is to handle the interaction with RFID reader. It can be any kind of an interface.

#### Data storage and processing

The data captured by the RFID reader is raw or unprocessed data. The function of the event manager is to process the data and make it useful. This is achieved by filtering, aggregation and transformation[15] of the captured data.

#### Application interface

This is the top layer of the RFID middleware. The primary objective of this component is to provide [14] a standard output that allows an application to receive filtered data from the RFID reader. In addition, the application-level interface should enable the end user to manage and improvise the data from the reader so as to provide multi-functionality.

All three components that make up an RFID system have been described in this section. The summary of the relationship between each of these components and what they ultimately aim to achieve is shown in Figure 1. The raw data from the RFID tag is captured by the RFID reader and send to the RFID middleware, which processes this data and outputs it to an end user computer. The data which is available at the computer can be applied in any way befitting the end user.



Fig 1: A typical RFID system

#### V. RFID HARDWRE DESIGN

Section IV described the basic components of an RFID system. This section will explain how the hardware

components of the RFID system mentioned in Section IV are designed and constructed. The hardware section consists mainly of the RFID reader and ensuring that it is able to collect data from the RFID tags.

The students are provided with a student card on registration. These cards are themselves RFID tags conforming to the ISO 14443 B standard for proximity cards which worked on the high frequency range at 13.56MHz.

These cards have an internal antenna and a unique identification number (UIN) programmed into them. The unique identification number is not known to user, in this case the student; rather distinction between student cards is made using the actual student number assigned to the student.

The RFID reader that needs to be designed should be capable of reading the high frequency RFID tags. The RFID reader as mentioned in Section IV needs to have HF interface module and a control unit. The ACG HF Multi ISO reader module was used as the control unit.

The ACG HF Multi ISO reader module, shown in Figure 2, is conformed to read RFID tags that adhere to the ISO 14443A standard. A serial port interface was created to link the control unit with the RFID middleware. Figure 3 shows the schematic of this connection.



Fig 2 : The ACG HF Multi ISO reader module.



Fig 3: Schematic showing the connection of reader module to the RFID middleware  $% \left( {{{\rm{B}}_{{\rm{B}}}} \right)$ 

The HF interface which is used to capture data from the RFID tag was achieved by designing and constructing an

antenna which operates at 13.56MHz. The circuit diagram used in the design of the antenna is shown in Figure 4.



Fig 4: Schematic for HF interface antenna

#### VI.RFID MIDDLEWARE DESIGN

The RFID middleware is the part that needs to be designed next. Section IV details the three components of the RFID middleware as the RFID reader adapter, data storage and application level interface. The purpose of the RFID middleware is to process the raw data from the RFID tag. To achieve this, the RFID reader unit needs to be connected to an end user computer.

The RFID reader is interfaced to the end user computer using an RS232 cable; hence data collected from the RFID tags reaches the serial port of the end user computer. This is where the data needs to be stored, filtered and transformed so that it can be used for registering the attendance of students.

It is to store, filter and transform the data received via the serial port that Java and the Apache Derby are used. Apache Derby[16] is an open source relational database implemented entirely in Java. The advantage of using the Apache Derby is that it is based on Structured Query Language (SQL), Java Data Base Connector (JDBC) and Java.

The relationship between SQL, JDBC and Java in relation to this paper is described in this section. This is done by first looking at database program, MySQL, used for storing data from the RFID tags. Secondly it shows how Java program is used transform the data such that it can be used to automate the attendance register. JDBC is the link between Java and the MySQL database.

#### MySQL database

MySQL is a Relational Database Management System (RDBMS) that uses SQL. It is open source, which is the major attraction behind using it. Data is stored in MySQL in database objects called tables. A table is a collection of related database entries and it consists of rows and columns. A database may contain one or more tables. Each table is identified with a name.

A database table is created using certain SQL statements. There are basically two types of SQL statements: the Data Manipulation Language (DML) and the Data Definition Language (DDL). The Data Manipulation Language statements are used for querying and updating a database table. Some of the DML commands are given below.

- SELECT command used for extracting data from a database
- UPDATE command used for updating data in a database
- DELETE command used for deleting data from a database
- INSERT INTO command used for inserting data into a database.

The Data Definition Language statements are used for creating and deleting database tables. They also provide indexes, specify links between tables and impose constraints between tables. Some of the commonly used DDL statements are as follows:

- CREATE DATABASE used for creating new database tables
- ALTER DATABASE used for modifying a database
- CREATE TABLE used for creating a new table
- ALTER TABLE modifies a created database table
- DROP TABLE deletes a table
- CREATE INDEX creates a search key for database tables
- DROP INDEX deletes a search key for database table.

In order to automate the attendance register, there needs to be two database tables. The first database table, referred to as ADMIN, contains three columns. The first is the unique identification number of the card. The second column contains the initials and surname of the student and lastly the student number. This table is created manually. The screen shot of the ADMIN table is shown Figure 5.

The second database table that needs to be created must be updated automatically as the students stream into the class and scan their RFID tagged student cards by the RFID reader which is connected to a laptop. The automatic data entry is done using JAVA and JDBC.

#### Java programming and JDBC

The data scanned from the RFID tags are brought to the serial port of the computer. A program is written in Java to retrieve the data from the serial port and bring it to the Java platform. The scanned data corresponds to the unique identification number (UIN) of the RFID tag which is in the ADMIN table.

If there is a match between the between the two numbers the data is entered with the student number and a date stamp to the second table, referred to here as the ATTENDEE table. This is shown in Figure 6.

MySQL Command Line Client					
Enter password: ******* Welcome to the MySQL monitor. Commands end with ; or \g.					
Server vers	ion: 5.0.51a-commun	ity-nt MySQL Com	munity Edition (GPL)		
Type 'help;	' or '\h' for help.	Type '\c' to cl	ear the buffer.		
asu (Insum	attendance:				
Database ch	anged				
mysql≻ sele	ct * from admo;				
ERROR 1146	(42802): Table 'att	endance.admo' do	esn't exist		
mysql≻ sele	ct * from admin;				
CARDNO	STUDENTNAME	STUDENTNUMBER	-+ 		
22458801	F W LOXTON	206039883			
22458803	I M R RAMAFOLE	208004530			
22458802	¦ MAKHETHA K E	207058385	1		
22458804	M J MOSITO	208008543			
22458805	M O MOKCHACHANE	208040269			
1 22458805	I N MURUBHNE	208054201			
i 22458806	I P P MPELO	i 208055517 · 2000EE707			
22450007	I M M MORHETHO	2000000777			
22458809	P NTHSHIDI	208080074			
22458810	D H LEMPE	208080538			
22458811	I M C THAELE	208080570			
22458812	I N MUJOVO	208080678			
22458813	I V SENEKAL	208080872			
22458814	S MOTSOANE	208082476			
i 22458815	I G PRINCE				
i 22458816 ! 99750019	I G MODUPI	1 207018704			
22458818	! X ASSEGAT	207022043			
22458819	S F SAMBO	209024992			
22458819	B M NDABA	209027347			
22458820	P E SESELE	l 209034122	1		
22458821	M K VAN RENSBURG	209037369			
22458822	I M M MPHRIME	209047925			
1 22458822	I M M MPHRIME	209047925			
i 22458823	I K K MHKIHINUS	i 207050272			
22450024	I I N NIHAHISHNE	207034223			
22458826	M MOTAHANE	209056339			
22458827	I N PULENYANE	209062851			
22458828	I T NTLHOKOE	209062878			
22458829	L MOKALANYANE	209070099			
22458830	S L MDHULI	209075376			
i 22458831	I M P SEHULHRU	209075813			
1 22400002 1 22400002	I H H DOISHELO	1 207070177			
+	+	+	+		
36 rows in set (0.02 sec)					
nysql> 🗖					

Fig 5: ADMIN database table

MySQL Command Line Client			
206039883   F	W LOXTON	2008-01-21	
207058385   MA	KHETHA K E	2008-01-21	
208008543   MC	SITO M J	2008-01-21	
208054201   MU	RUBHNE N	2008-01-21	
208055517 i P	P NPELO i	2008-01-21	
208055797 ! 1	C MOKHELE	2008-01-21	
208080074 P	I NTHSHIDI	2008-01-21	
208080538   D	H LEMPE	2008-01-21	
208080570   M	C THAELE	2008-01-21	
208080678 I N	MUJOUO	2008-01-21	
208080678 i N	CENEVOI I	2008-01-21	
208081623 ! G	PRINCE !	2008-01-21	
208082476 S	MOTSOANE	2008-01-21	
209018704   T	G MODUPI	2008-01-21	
209022043   J	F_MALAN_	2008-01-21	
209022345   8	ASSEGAAI	2008-01-21	
207024772 i S	E SHMBU i	2008-01-21	
207027347 I D 200024122 I D	F CECTE I	2008-01-21	
209037369 M	K UAN RENSBURG	2008-01-21	
209054743   M	GUMBI	2008-01-21	
209056339   M	MOTAHANE I	2008-01-21	
209070099 L	N MOKALANYANE	2008-01-21	
207078513 i M	P SEHULHRU i	2008-01-21	
207070177 I H	MIKONI !	2008-01-21	
209027347 B	M NDABA	2008-01-28	
209090316   B	MUKANI	2008-01-28	
209090197   A	A BOTSHELO	2008-01-28	
209078513 ; M	P SEHULARO	2008-01-28	
207070077 I L 2000EC220 I M	MOTOBONE !	2008-01-28	
209054743 M	GUMBI	2008-01-28	
209037369   M	K VAN RENSBURG	2008-01-28	
209034122   P	E SESELE	2008-01-28	
209024992   S	E SAMBO I	2008-01-28	
209022043 ! .	F MALAN !	2008-01-28	
208082476 S	MOTSOANE	2008-01-28	
208080872 ¦ V	SENEKAL I	2008-01-28	
208080678   N	MUJOVO	2008-01-28	
208081623   G	C THOFIE I	2008-01-28	
208080074 ! P	I NTHSHIDI !	2008-01-28	
208055797   K	C MOKHELE	2008-01-28	
208080031   M	M MAKHETHA I	2008-01-28	
208055797   KC	MOKHELE	2008-01-28	
209052851 i N	PULENYHNE I	2008-01-28	
206039883 ! F	U LOXTON !	2008-02-04	
206039883   F	W LOXTON I	2008-01-28	
208004530   M	R RAMAFOLE	2008-02-04	
208055517   P	P MPELO	2008-02-04	
208055797 1 К	M MOKHETHO	2008-02-04	
208080570 M	C THAELE	2008-02-04	
208081623   G	PRINCE	2008-02-04	
209022043 J	F MALAN	2008-02-04	
209022345   X	ASSEGAI I	2008-02-04	
207027347 i B 209024122 i D	E SESELE	2008-02-04	
209050292	R MARTHINUS	2008-02-04	
209054743 M	GUMBI	2008-02-04	
209062878   T	NTLHOKOE	2008-02-04	
209070099 L	N MOKALANYANE	2008-02-04	

Fig 6: ATTENDEE database table

#### VII.RESULTS AND DISCUSSION

This application was developed and piloted with the intention to automate one of the most critical aspects of an educational institution being attendance monitoring. The most crucial element which was wasted when the procedure was manual was time.

With this procedure the attendance was registered within a matter of seconds. The second aspect was categorizing the attendance to specific dates and tabulating the total attendance for each student. This was also possible with the MySQL software. This was tested during the pilot run of four classes.

From a technical point of view, in the literature review in section IV, the range of a HF RFID tag was said to be about 20cm. In actual practice this was much shorter. On average it was between 3 and 5cm. A picture representing the test is shown in Figure 7.



Fig 7: Testing the average range of the HF RFID tag

The challenges posed by anti-collision is negated by algorithms written into the reader module while duplication of student registration is countered using normalization techniques used in the MySQL database

#### VIII.CONCLUSIONS

The purpose of this paper was to design a system that would automate the student attendance register. After looking into the different technologies available and the existing resources in the university, RFID technology was deemed fit for this purpose. RFID technology on its own will not be able to process the raw data stored in the RFID tags. Hence MySQL and Java programming were combined to store and transform the scanned data from the RFID tags respectively to create an attendance register.

The pilot run showed that most challenges that were explained in the introduction with the manual attendance monitoring like time consumption, data loss, error proneness and difficulty in tabulation were overcome using this system.

Questions can be asked as to authenticity of the card holder, but the argument would be that monitoring the automatic attendance register is much easier than the manual system.

#### REFERENCES

[1] J. Fantuzzo and S. Grim, "An evaluation of a community-wide school based intervention to reduce truancy," *Psychol. Sch.*, vol. 42, no. 6, pp. 657–667, 2005.

[2] M. Landin and J. Perez, "Class attendance and academic achievement of pharmacy students in a European University," *Curr. Pharm. Teach. Learn.*, vol. 7, no. 1, pp. 78–83, 2015.

[3] L. Stanca, "The Effects of Attendance on Academic Performance: Panel Data Evidence for Introductory Microeconomics," *J. Econ. Educ.*, vol. 37, no. 3, pp. 251–266, 2006.

[4] D. Romer, "Do students go to class-Should they?," *J. Econ. Perspect.*, vol. 7, no. 3, pp. 167–174, 1993.

[5] H. Vermaak and RB Kuriakose, "Using reflective practices to reduce dropout rates among first year students at a University of Technology, a South African perspective," in *Frontiers in Education conference*, 2014.

[6] M. Castells, "Lecture on Higher Education," 2009. [Online]. Available: http://www.chet.org.za/files/uploads/events/Seminars/Higher Education and Economic Development.pdf?download=1. [Accessed: 20-May-2010].

[7] M. Ramakrishnan and R. Josphineleela, "An Efficient Automatic Attendance System using Fingerprint Verification Technique," *Int. J. Comput. Sci. Inf. Secur.*, vol. 10, no. 3, pp. 264–269, 2012.

[8] M. Ajana, E. Khaddar, M. Boulmalf, and H. Harroud, *RFID Middleware Design and Architecture*. 2011.

[9] K. Finkenzeller, *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards, Radio Frequency Identification and near-Field Communication*, 3rd ed. Wiley and Sons, 2010.

[10] S. A Weis, "RFID (Radio Frequency Identification): Principles and Applications," in *Emerging Tecnologies*, Harvard, 2010, pp. 1–23.

[11] H. Al-Mousawi, "Performance and reliability of Radio Frequency Identification (RFID) : theoretical evaluation and practical testing in relation to requirement from use in Abu Dhabi Sewerage Directorate," Universitetet i Agder, 2004.

[12] K. Michael and S. Wamba, "An information system design theory for and RFID," *Research Online*, 2008. [Online]. Available: http://ro.uow.edu.au/cgi/viewcontent.cgi?article=1720&context=infopape rs. [Accessed: 09-Oct-2015].

[13] J. V Gorabal and D. H. Manjaiah, "RFID Concepts, Applications and Issues," *International. Journal of. Engineering . Resources and. Technology.*, vol. 2, no. 12, pp. 3319–3321, 2013.

[14] F. Lin, B. Chen, C. Y. Chan, C. H. Wu, W. H. Ip, A. Mai, H. Wang, and W. Liu, "The Design of a Lightweight RFID Middleware," *Internation. Journal of Engineering and. Business Management.*, vol. 1, no. 2, p. 1, 2009.

[15] C. Floerkemeier and M. Lampe, "RFID middleware design: addressing application requirements and RFID constraints," in *Proceedings of the 2005 Joint Conference on Smart Objects and Ambient Intelligence: Innovative Context-aware Services: Usages and Technologies*, 2005, no. october, pp. 219–224.

[16] "Getting Started with Apache Derby," 2012. [Online]. Available: http://db.apache.org/derby/docs/10.3/getstart/getstartderby.pdf. [Accessed: 10-Oct-2015].

#### Normal Distribution Transform Graph-based Point Cloud Segmentation

William R. Green<sup>1</sup> and Hans Grobler<sup>2</sup>

Abstract—We present a graph-based algorithm for segmenting point cloud scenes using criteria based on the combination of spatial, geometric, and appearance features. An octree data structure is employed to organize the point cloud data. The voxel space is used to create a Normal Distribution Transform Feature Representation (NDT-FR) to model the underlying sensor data and corresponding features in a probabilistic manner. The proposed segmentation algorithm uses the Hellinger distance calculated on local statistics stored in neighboring voxels to define the edge weights of the graph. Rather than choosing a specific feature for edge weight calculation, our approach has the ability to combine multiple features into a single edge weight without the need to find an appropriate normalization scheme. We verify our algorithm on multiple indoor scenes and perform a qualitative evaluation. We also show how our edge weighting scheme can increase the accuracy of object boundaries in the final segmentation.

#### I. INTRODUCTION

Low-cost depth sensors and full 3D laser range finders are widely available and have the ability to produce dense 3D point clouds. The geometric and spatial properties of objects present in the captured scene are implicitly embedded in the point cloud. The ability to extract these properties can greatly assist in solving problems such as object detection and activity recognition. Segmentation is a crucial step in many vision algorithms where some form of scene understanding is required. Segmentation techniques aim to cluster points with similar properties into multiple homogeneous regions based on some proximity and/or feature distribution criterion. The coarser collections of data produced by the segmentation algorithm should be perceptually important and support the global features of the captured scene.

The accuracy of the segmentation directly depends on the quality and structure of the input data. Due to the limitations of 3D sensors, point cloud data are usually unevenly sampled and noisy [1]. Furthermore, the unstructured nature of a point cloud poses several challenges when performing spatial and neighborhood operations. Data structures such as an octree have been used widely in literature to create a structured voxelized representation of a point cloud [2]. The nature of 3D segmentation algorithms allows for efficient use of the discretized and organized representation when considering neighborhood and data query operations [3]–[5].

Due to the sheer amount of data present in a point cloud, traditional segmentation techniques such as region growing or split-and-merge algorithms are too computationally expensive for general use [6]. Felzenszwalb and Huttenlocher

<sup>1</sup>William R. Green and <sup>2</sup>Hans Grobler are with the Department of Electrical, Electronic and Computer Engineering, University of Pretoria, South Africa. robertgreen@gigacomp.co.za, hans.grobler@up.ac.za

introduced a graph-based segmentation algorithm for 2D color images that scales almost linearly with the number of pixels in the image [7]. More recently the algorithm has been extended to the segmentation of color 3D laser point clouds [8]. The extended graph-based segmentation algorithm has also been used successfully as a pre-processing step for object detection in cluttered scenes [9], [10]. Graph-based segmentation algorithms require the creation of edges between each set of connected vertices. Edges are defined by calculating a weight that represents the correspondence between the two connected vertices. The criteria and features chosen to calculate the edge weight have a direct influence on the stability and accuracy of the segmentation algorithm.

In this paper we present a new graph-based segmentation algorithm for segmenting 3D point cloud data. We introduce the concept of a Normal Distribution Transform Feature Representation (NDT-FR) that is based on an octree data structure. The NDT-FR can model local statistics on many different features extracted from the point cloud. A key contribution is the use of a new edge weighting scheme that is based on the difference criterion between local statistics stored in neighboring voxels. We evaluate our segmentation algorithm by performing a qualitative analysis on the ability to distinguish between objects in the segmented scene.

#### II. BACKGROUND

Our proposed segmentation algorithm is based on the Felzenszwalb segmenter [7]. We extend the algorithm by using an octree to create a 3D voxel grid of the point cloud and build the graph using neighboring cells in the voxel space. We model the point and feature distribution in each voxel using the Normal Distribution Transform (NDT) [11]. We further extend the standard Felzenszwalb segmenter by implementing a statistical distance measure between the feature distributions to define our graph edge weights.

#### A. Data Structure and Representation

Range measurements representing a 3D environment can be modeled using a grid of 3D volumes that are commonly referred to as voxels or cells, and was first introduced by Moravec and Elfes [12]. The voxel space is usually managed in an octree data structure allowing for efficient data storage, queries, and manipulation operations. Each node in the octree represents a voxel or cubic volume in 3D space. Information such as occupancy can be stored in each voxel [2].

The NDT representation is very similar to the occupancy grid approach but models the observed range measurements within each voxel with a normal distribution [11]. Each cell in the 3D grid is associated with a sampled mean **q** and covariance  $\Sigma$  representing the multivariate normal distribution  $\mathcal{N}(\mathbf{q}, \Sigma)$ . NDT-Maps have been used in 3D mapping algorithms and can function in dynamic environments [13], [14]. For each acquired sensor measurement the corresponding voxel in the NDT-Map needs to be updated. The update operation entails computing the new sampled mean and covariance for a voxel. The update operation proposed by Biber *et al.* [11] requires that all points observed in a voxel should be stored. The approach is not memory efficient nor computationally scalable. Saarinen *et al.* [13] detail a method for recursively estimating the mean and covariance of a multivariate normal distribution. Given two sets of measurements  $\{x_k\}_{k=1}^m$  and  $\{y_k\}_{k=1}^n$ , the recursive update rules for the mean and covariance take the form

$$\mathbf{q}_{x\oplus y} = \frac{1}{m+n} \left( m\mathbf{q}_x + n\mathbf{q}_y \right) \tag{1}$$

and

$$\boldsymbol{\Sigma}_{x \oplus y} = \frac{1}{m+n-1} \left( \boldsymbol{\Sigma}_x + \boldsymbol{\Sigma}_y + \frac{nm}{n+m} \mathbf{u} \mathbf{u}^T \right), \quad (2)$$

where  $\mathbf{u} = \mathbf{q}_x - \mathbf{q}_y$ . For the update rules (1) and (2), it is only necessary to store the number of samples together with the current sampled mean and covariance. In our segmentation algorithm we explore the usage of the NDT representation to model other features such surface normals and color.

#### B. Graph-based Segmentation

Given the graph G = (V, E) with vertices  $V = [v_1, ..., v_i]$ that are data points from a joint point cloud and E a set of edges. The aim is to segment the graph into a set of regions  $S = [s_1, ..., s_i]$  so that the union of these segments produce the joint point cloud V. In the simplest case, each point  $d_i$  in the cloud represents a vertex  $v_i$ . An edge  $e_{ij} \in E$  is defined as  $edge(v_i, v_j)$  and has a corresponding edge weight  $w_{ij}$  that defines the dissimilarity between the vertices  $v_i$  and  $v_j$ .

We consider the Felzenszwalb segmenter as the basis for our graph-based segmentation technique [7]. The algorithm segments an undirected graph by considering the internal variation of properties in neighboring regions. For the case of a 2D RGB image the algorithm creates a graph where each pixel is a vertex and the graph edges are created by considering a 8-connected neighborhood in the image grid. Edge weights are determined by calculating the color difference between the pixels corresponding to the vertices of the edge.

The first step of the algorithm is to sort the edges Ein increasing order according to  $w_{ij}$ . Each graph segment is defined as  $s_i = \{d_{s_i}, thresh_{s_i}\}$ , where  $d_{s_i}$  represents all data points belonging to the segment and  $thresh_{s_i}$  is defined as the merging threshold of the region. The graph is segmented by processing the edges in order and merging the segments  $s_i$  and  $s_j$  corresponding to the edge vertices only if the edge weight  $w_{ij}$  is less than  $thresh_{s_i}$  and  $thresh_{s_j}$ . If the merging criteria is met, the new segment's threshold  $thresh_{s_{ij}}$  can be calculated as

$$thresh_{s_{ij}} = w_{ij} + \frac{k}{|s_{ij}|}.$$
(3)

The Felzenszwalb segmenter only relies on the two parameters T and k that are typically empirically chosen based on the type of graph weights used and the type of scene segmented. Each segment's merging threshold is initialized with T. The parameter k in (3) directly affects the granularity of the final segmentation. During the initial phase of the algorithm where the edge weights  $w_{ik}$  are very small, the term  $\frac{k}{|s_{ij}|}$  in (3) dominates the new threshold value. As the segment size stabilizes and the weights become larger the threshold  $thresh_{s_{ij}}$  approaches the largest intra-segment weight in  $s_i$  and  $s_j$ .

The image grid based Felzenszwalb segmenter has been extended successfully for the segmentation of 3D data [8]. Finman *et al.* [9], [15] defined a graph where the vertices were points in cloud slices produced by the Kintinuous RGB-D mapping system [16]. A surface mesh as the basis for the graph can also be used to ensure that objects in the foreground are not connected to objects in the background when the graph edges are created [8], [10]. For efficiency and fast neighborhood queries we make use of the octree data structure and a voxelized point cloud to implicitly represent the graph needed for the Felzenszwalb segmenter. Graph edges are constructed by performing neighborhood operations in the octree and vertex information is stored in voxels corresponding to the nodes in the octree.

The criteria and features used to calculate the edge weights directly influence the final segmentation and the accuracy of the extracted object boundaries. A simple approach is to implement the color distance in Euclidean space [8]. Surface normals can also be used as features for the creating graph edges. Both an unbiased and biased surface normal weighting scheme have been proposed in literature. For the unbiased case the simple angle difference between surface normals is used [8]. The biased normal weighting scheme on the other hand is a curvature aware metric that favors convex edges and penalizes concave edges [17].

In many cases it would be advantageous to make use of multiple metrics to define a single edge weight. A white box on a white table can be separated easily when looking at the surface normal space. In the case of a white paper on a black desk, the color space would yield the necessary object boundaries. Simply combining two distance metrics is not an easy task since the metrics are usually in fundamentally different units. Furthermore, the first step after graph construction in the Felzenszwalb segmenter algorithm is sorting the edges based on their corresponding weight. When considering multiple edge weights the sorting process becomes unclear. Strom et al. [8] proposed choosing one of the weights to sort on and simply introduce another threshold when considering a join between regions. Sorting on a single weight still has the effect of biasing the segmentation results for a specific distance metric. In this work we propose a new edge weighting scheme that has the ability to combine multiple edge weights by considering a unit-less statistical distance metric.

#### C. Distance Metrics for Feature Comparison

Given two feature histograms it is often necessary to determine the histogram distance. Several distance metrics exist in literature to determine the dissimilarity between feature histograms [18], [19]. The  $\chi^2$  distance metric has been used to compare CIE LAB histograms to create edge weights for graph-based segmentation [20], [21]. Our algorithm employs the NDT representation as the basis for features resulting in continuous feature distributions rather than discrete histograms. For an undirected graph it is necessary to use a symmetric difference measure to calculate edge weights. Metrics such as the Kullback-Leibler divergence and  $\chi^2$  distance that can be used to calculate the distance between probability measures are not symmetric. The Hellinger distance selected for use in this approach is one of the statistical distance metrics that is symmetrical. Furthermore, it is bounded in the range [0, 1] and takes into account the uncertainty of the feature distributions [22].

Using the Rényi divergence it is possible to derive the Hellinger distance between two multivariate normal distributions [23]. The Rényi divergence between two multivariate normal distributions  $\mathcal{N}_1(\mathbf{q}_1, \boldsymbol{\Sigma}_1)$  and  $\mathcal{N}_2(\mathbf{q}_2, \boldsymbol{\Sigma}_2)$  for order a is given by

$$D_a^1(\mathcal{N}_1, \mathcal{N}_2) = \frac{1}{a-1} \log \int\limits_{\chi} K_a^*(\mathcal{N}_1, \mathcal{N}_2), \tag{4}$$

where

$$K_a^*(\mathcal{N}_1, \mathcal{N}_2) = \int_{\chi} f_{\mathcal{N}_1}(\mathbf{x})^a f_{\mathcal{N}_2}(\mathbf{x})^{1-a} du(\mathbf{x}).$$
(5)

The relationship between the squared Hellinger distance  $H^2(\mathcal{N}_1,\mathcal{N}_2)$  and  $K^*_a(\mathcal{N}_1,\mathcal{N}_2)$  can be shown to be

$$H^{2}(\mathcal{N}_{1},\mathcal{N}_{2}) = 1 - K^{*}_{\frac{1}{2}}(\mathcal{N}_{1},\mathcal{N}_{2}).$$
 (6)

The expression for  $K_a^*(\mathcal{N}_1, \mathcal{N}_2)$  takes the form

$$K_a^*(\mathcal{N}_1, \mathcal{N}_2) = \frac{\det\left(a\boldsymbol{\Sigma}_2 + (1-a)\boldsymbol{\Sigma}_1\right)^{-\frac{1}{2}}}{\det\left(\boldsymbol{\Sigma}_1\right)^{\frac{a-1}{2}}\det\left(\boldsymbol{\Sigma}_2\right)^{-\frac{a}{2}}} \times \exp\left\{\frac{a(a-1)}{2}\mathbf{u}^T\left(a\boldsymbol{\Sigma}_2 + (1-a)\boldsymbol{\Sigma}_1\right)^{-1}\mathbf{u}\right\}.$$
 (7)

From (7) the squared Hellinger distance can be calculated as

$$H^{2}(\mathcal{N}_{1},\mathcal{N}_{2}) = 1 - \frac{|\boldsymbol{\Sigma}_{1}|^{\frac{1}{4}}|\boldsymbol{\Sigma}_{2}|^{\frac{1}{4}}}{\left|\frac{1}{2}\boldsymbol{\Sigma}_{1} + \frac{1}{2}\boldsymbol{\Sigma}_{2}\right|^{\frac{1}{2}}} \times \exp\left\{-\frac{1}{8}\mathbf{u}^{T}\left(\frac{1}{2}\boldsymbol{\Sigma}_{2} + \frac{1}{2}\boldsymbol{\Sigma}_{1}\right)^{-1}\mathbf{u}\right\}. (8)$$

In (7) and (8)  $u = q_1 - q_2$ .

### III. NORMAL DISTRIBUTION TRANSFORM SEGMENTATION

In this section we present our graph-based segmentation algorithm. We introduce the idea of a Normal Distribution Transform Feature Representation (NDT-FR) and show how local statistics stored in voxels can be used for a probabilistic graph edge weighting scheme. The algorithm first integrates the sensor readings and corresponding features into the NDT-FR using the NDT recursive update rules. We construct a graph from the NDT-FR where each vertex in the graph initially represents a voxel in the discretized environment. Neighboring voxels are used to create graph edges based on spatial, geometric, and appearance features stored in the NDT-FR. Finally, the graph is segmented using the Felzenszwalb segmenter.

#### A. Data Structure for Feature Enriched 3D Data

The standard NDT-Map approach primarily integrates range measurements into the 3D voxel grid. In many cases each point in a cloud can have a corresponding surface normal and appearance features. We propose a NDT-FR where these features are also integrated into the voxel space using the NDT update rules. The NDT-FR contains for every voxel a set of normal distributions modeling the range measurements and the corresponding feature spaces.

Although arbitrary features are supported, in this version of the implementation we only consider point position, surface normal, and color features. As input to our segmentation algorithm we use a RGB-D frame. A point cloud can be obtained by simply using the camera parameters and the depth image. The color feature for every point is determined by looking up the color in the RGB image for each point in the depth image. Following the standard approach we smooth the RGB image with a  $\sigma = 0.8$  Gaussian filter [7]. The color and depth images should be registered for accurate data association. A surface normal for each point in the cloud is computed by exploring the structured nature of the RGB-D data [24]. We consider pixel neighborhoods rather than costly spatial neighborhoods for surface normal estimation as proposed by Lee *et al.* [25].

Due to the limitations of Time-of-Flight sensors the depth maps can be very noisy [1]. Segmentation algorithms that rely on feature extraction such as normal estimation employ noise removal algorithms as a pre-processing step [26]. We implement a bilateral filter to reduce the amount of noise in the raw depth map I [27]. The filtered depth image  $I_f$  is obtained by applying the filter

$$I_{f} = \frac{1}{w_{p}} \sum_{q \in S(p)} G_{\sigma_{s}} \big( \|p - q\| \big) G_{\sigma_{r}} \big( \|I(p) - I(q)\| \big) I(q),$$
(9)

where  $w_p$  is the normalization factor and  $G_{\sigma_s}$  and  $G_{\sigma_r}$  are Gaussian kernels  $G_{\sigma}(x) = \frac{1}{2\pi\sigma} \exp\left(-x^2/(2\sigma^2)\right)$ . The filtered point cloud together with the extracted fea-

The filtered point cloud together with the extracted features are used to create the NDT-FR. We implement an octree as the underlying data structure to discretize the point cloud into a 3D voxel grid. Each cell  $c_i$  in the voxel space is represented by the set of parameters

$$c_i = \{\mathbf{q}_p, \mathbf{\Sigma}_p, \mathbf{q}_n, \mathbf{\Sigma}_n, \mathbf{q}_c, \mathbf{\Sigma}_c, N\},\tag{10}$$

where  $\mathbf{q}_p$ ,  $\Sigma_p$ ,  $\mathbf{q}_n$ ,  $\Sigma_n$ ,  $\mathbf{q}_c$ , and  $\Sigma_c$  are the mean and covariance parameters for the normal distributions representing the point position, surface normal, and color distribution in a voxel. The parameter N corresponds to the number of points



(a) Voxel surface normal features. (b) Voxel color features.

Fig. 1: Two feature spaces integrated into the NDT-FR.

Algorithm 1 Build Graph			
Input: Tree: Octree			
Output: E: Graph edges			
1: $E \leftarrow \emptyset$			
2: for all $c_i \in Tree$ do			
3: $N_{c_i} \leftarrow \left\{ c_j \in Tree \mid \ c_i - c_j\  < r \right\}$			
4: for $c_j \in N_{c_i}$ do			
5: $w_{ij} \leftarrow \sum_{k=1}^{k=K} a_k H^2 \big( \mathcal{N}_i(\mathbf{q}_k, \mathbf{\Sigma}_k), \mathcal{N}_j(\mathbf{q}_k, \mathbf{\Sigma}_k) \big)$			
6: $E \leftarrow \{E \cup edge(c_i, c_j)\}$			
7: end for			
8: end for			
9: Sort $E$ in ascending order according to $w_{ij}$			

used to obtain the normal distribution parameters in a voxel so far. For each point in the cloud the position of the voxel that was "hit" is determined. The voxel parameters (10) are updated with the NDT update rules (1) and (2), and the parameter N is increased. In Fig. 1 we visualize the surface normal and color feature space integrated in a NDT-FR.

#### B. Graph Construction and Segmentation

The underlying octree data structure of the NDT-FR can implicitly be used to create the graph for segmentation. Algorithm 1 describes our method for constructing the graph and takes as input the tree representing the NDT-FR created from the data points and features extracted from a single scene point cloud. For each voxel  $c_i$  in the tree the neighboring cells  $N_{c_i}$  within a radius r are obtained. In our implementation we set r to be twice the resolution of the octree to restrict the number of edges in the graph. In line 5 we construct an edge  $edge(c_i, c_j)$  between the voxel  $c_i$  and one of its neighbors  $c_j$  in  $N_{c_i}$  using the Hellinger distance (8). K is the number of features stored in each cell  $c_i$  and  $\mathcal{N}_i(\mathbf{q}_k, \boldsymbol{\Sigma}_k)$  is the normal distribution for feature k in  $c_i$  (10). The bounded nature of the distance metric allows us to calculated the final edge weight by simply taking a weighted average of all the distances between the respective feature distributions. Each edge now only has a single weight encapsulating the difference between multiple feature spaces. The final step is to sort the edges in ascending order according to the calculated edge weights  $w_{ii}$ .

The Felzenszwalb segmenter is implemented to segment the graph as described in Algorithm 2. The segmentation S is initialized by assigning to each voxel an unique segment and

Algorithm 2 Graph-segmentation **Input:** D: Structured Point Cloud **Output:** S: Set of segments 1:  $Tree \leftarrow \text{Build}_N\text{DT}_FR(D)$ 2:  $E \leftarrow \text{Build}_\text{Graph}(Tree)$ 3:  $S \leftarrow \emptyset$ 4: for i = 1...numVoxels do 5:  $s_i \leftarrow \{c_i, T\}$ 6: end for for all  $e_{ij} \in E$  do 7: if  $w_{ij} < \min(thresh_{s_i}, thresh_{s_j})$  then 8:  $\dot{c_{s_m}} \leftarrow c_{s_i} \cup c_{s_j} \\ thresh_{s_m} \leftarrow w_{ij} + \frac{k}{|c_{s_m}|}$ 9: 10:  $s_m \leftarrow \{c_{s_m}, thresh_{s_m}\} \\ S \leftarrow \{S \setminus \{s_i \cup s_j\}\} \cup s_m$ 11: 12: end if 13: 14: end for

initial merging threshold T. We implement a disjoint-set data structure with union by rank and path compression to keep track of the segments and the corresponding voxels belonging to each segment [28]. The algorithm proceeds by processing all edges in the graph in increasing order. If the merge criterion in line 8 is met, the segments connected by the edge are merged with a union operation. The segmentation S is updated with the new segment  $s_m$  having a merging threshold calculated according to (3). Since very small regions most likely do not correspond with an object we perform a last iteration over all edges and greedily merge any segments that are smaller than some pre-defined threshold minSize. A find operation on each of the initial segments in the disjoint-set is then performed and the corresponding voxels in the tree is updated with the correct segmentation labels.

#### **IV. RESULTS**

We tested our segmentation algorithm on various sequences from the TUM RGB-D benchmark [29], and the ICL-NUIM dataset [30]. The TUM dataset is used since there are several noise characteristics present in the RGB-D frames whereas the ICL-NUIM dataset is a synthetic noisefree RGB-D sequence. An octree with a 0.008 m resolution was used in all of our experiments. The parameters T and k were kept constant for our segmentation algorithm and were chosen empirically to be 0.3 and 12.0 respectively. To eliminate very small segments we set the minSize parameter to 90 voxels. Furthermore, we only consider points in the cloud with a distance less than 3.5 m when using the TUM dataset to reduce sensor noise artifacts. In the first experiment we compare our segmentation approach to the algorithm proposed by Strom et al. [8]. More specifically, we evaluate the working of our new edge weighting scheme by qualitatively comparing segmentations produced by the two approaches and analyzing the effect of using multiple edge weights on the segmentation parameters. In a second experiment we show segmentation results for noise free and noisy RGB-D scenes. No timing results are presented since the current implementation was done in Python.





(a) Color only Hellinger distance weights.

(b) Surface normal only Hellinger distance weights.



(c) Combining surface normal and color Hellinger distance weights.

#### Fig. 2: NDT-FR edge weight combination.

Our segmentation results for a cluttered table scene in the TUM dataset is shown in Fig. 2. We examine the segmentation results using either surface normal or color difference to define edge weights as shown in Fig. 2a and Fig. 2b respectively. Fig. 2a and Fig. 2b are projected feature images from the NDT-FR. When using only color, an object such as the white cup in the lower right of the scene cannot readily be distinguished from the white table. Similarly, using only surface normals does not allow for accurate boundary detection between the blue book and telephone on the right side of the table. By combining the surface normal and color edge weights a superior segmentation can be obtained as shown in Fig. 2c.

Fig. 3 shows the segmentation results obtained from using the edge weighting scheme implemented by Strom *et al.* [8]. The segmentation when using only color difference in Euclidean space is shown in Fig. 3a whereas the segmentation obtained using only surface normal angle difference is shown in Fig. 3b. Combining the surface normal and color edge weights by sorting according to one weight and using both for the merging criteria is shown in Fig. 3c and Fig. 3d. Table I details the parameters used for the segmentation. The *minSize* parameter was once again set to 90 and the  $T_{rgb}$ and  $T_{norm}$  parameters were set to 100 and 0.8 respectively. Note the huge difference in parameter settings when using different features. Furthermore, the parameters have to be adjusted when the two edge weights are combined.

Comparing the results showed in Fig. 2 and Fig. 3 several observations can be made. Our probabilistic edge weighting scheme allows our algorithm to use the same set of segmentation parameters when using different features or combining features to define edge weights. When combining edge weights our algorithm does not bias the segmentation towards any one of the individual edge weights. From Fig. 3c and Fig. 3d it is clear that the segmentation is biased

TABLE I: Segmentation parameters for Fig. 3

Figure	3a	3b	3c, 3d
k <sub>rgb</sub>	2200	n/a	25000
$k_{norm}$	n/a	3.5	12.0





(a) Color only Euclidean distance weights.



(b) Surface normal angle only difference weights.



(c) Sorting according to color (d) Sorting according to surface weights.

Fig. 3: Edge weight combination by sorting [8].

according to the edge weight chosen to sort on.

We show our segmentation results for other scenes from the TUM and ICL-NUIM dataset in Fig. 4. In Fig. 4c and Fig. 4f we show segmentations obtained for two scenes from the ICL-NUIM dataset using surface normal and color features. The noise free dataset allows for the extraction of very accurate object boundaries. The final segmentation shown in Fig. 4i was obtained for a RGB-D frame in the noisy TUM sequence. For the segmentation we added the point position distribution in each voxel as a feature to be used for edge weight construction together with the surface normal and color features. The edge weights were also biased towards the color and normal features. Even for a noisy depth map and RGB image our segmentation algorithm accurately identifies most object boundaries.

#### V. CONCLUSION

In this paper we presented a graph-based segmentation algorithm that uses structured data from a RGB-D sensor. We introduced the concept of a NDT-FR that models the point distribution and corresponding features in a 3D point cloud as a set of normal distributions. Our algorithm uses a new edge weighting scheme that is defined by the statistical difference between two normal distributions. We showed that our proposed edge weighting scheme allows for combining features into a single edge weight without finding an appropriate normalization scheme. Furthermore, the segmentation parameters are not affected when new feature information is integrated into the final edge weights. The combination







Fig. 4: Segmentation results for different RGB-D scenes from the TUM and ICL-NUIM dataset. The left and middle columns show the input depth and RGB frames respectively and the right column depicts the final segmentation.

of features for defining edge weights allowed for a qualitative superior segmentation. However, depth image noise and inaccurate RGB and depth frame registration can cause erroneous segmentation boundaries. Future work will include a quantitative evaluation of our method with existing methods as well as incorporating contextual and other higher level information into our edge weighting scheme.

#### REFERENCES

- [1] J. Fu, S. Wang, Y. Lu, S. Li, and W. Zeng, "Kinect-Like Depth Denoising," in Proc. of IEEE Int. Symposium on Circuits and Systems (ISCAS), May 2012, pp. 512-515.
- [2] A. Hornung, K. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: A Probabilistic, Flexible, and Compact 3D Map Representation for Robotic Systems," Autonomous Robots, vol. 34, no. 3, pp. 189-206, 2013.
- [3] H. Woo, E. Kang, S. Wang, and K. Lee, "A new segmentation method for point cloud data," Int. Journal of Machine Tools and Manufacture, vol. 42, no. 2, pp. 167-178, 2002.
- [4] M. Wang and Y.-H. Tseng, "Incremental segmentation of lidar point clouds with an octree-structured voxel space," Photogrammetric Record, vol. 26, no. 133, pp. 32-57, 2011.
- [5] A. Vo, L. Truong-Hong, D. Laefer, and M. Bertolotto, "Octree-based region growing for point cloud segmentation," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 104, pp. 88-100, 2015.
- [6] A. Nguyen and B. Le, "3D Point Cloud Segmentation: A survey," in Proc. of IEEE Conf. on Robotics, Automation and Mechatronics (RAM), Nov 2013, pp. 225-230.
- [7] P. Felzenszwalb and D. Huttenlocher, "Efficient Graph-Based Image Segmentation," Int. Journal of Computer Vision, vol. 59, no. 2, pp. 167-181, 2004.
- [8] J. Strom, A. Richardson, and E. Olson, "Graph-based Segmentation for Colored 3D Laser Point Clouds," in *Proc. of IEEE/RSJ Int. Conf.* on Intelligent Robots and Systems (IROS), 2010, pp. 2131-2136.

- [9] R. Finman, T. Whelan, M. Kaess, and J. Leonard, "Toward Lifelong Object Segmentation from Change Detection in Dense RGB-D Maps, in Proc. of European Conf. on Mobile Robots (ECMR), 2013, pp. 178-185
- [10] A. Karpathy, S. Miller, and L. Fei-Fei, "Object Discovery in 3D scenes via Shape Analysis," in Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA), 2013, pp. 2088-2095.
- [11] P. Biber and W. Strasser, "The Normal Distributions Transform: A New Approach to Laser Scan Matching," in Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), vol. 3, Oct 2003, pp. 2743-2748
- [12] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA), vol. 2, Mar 1985, pp. 116-121.
- [13] J. Saarinen, H. Andreasson, T. Stoyanov, J. Ala-Luhtala, and A. Lilienthal, "Normal Distributions Transform Occupancy Maps: Application to Large-Scale Online 3D Mapping," in Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA), May 2013, pp. 2233-2238.
- [14] E. Einhorn and H.-M. Gross, "Generic 2D/3D SLAM with NDT Maps for Lifelong Application," in Proc. of European Conf. on Mobile *Robots (ECMR)*, Sept 2013, pp. 240–247.
- [15] R. Finman, T. Whelan, M. Kaess, and J. Leonard, "Efficient Incremental Map Segmentation in Dense RGB-D Maps," in Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA), May 2014, pp. 5488-5494.
- [16] T. Whelan, M. Kaess, M. Fallon, H. Johannsson, J. Leonard, and J. McDonald, "Kintinuous: Spatially Extended KinectFusion," in RSS Workshop on RGB-D: Advanced Reasoning with Depth Cameras, Sydney, Australia, Jul 2012.
- [17] F. Moosmann, O. Pink, and C. Stiller, "Segmentation of 3D Lidar Data in non-flat Urban Environments using a Local Convexity Criterion," in Proc. of IEEE Intelligent Vehicles Symposium, June 2009, pp. 215-220.
- [18] G. Hetzel, B. Leibe, P. Levi, and B. Schiele, "3D Object Recognition from Range Images using Local Feature Histograms," in Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR), vol. 2, 2001, pp. II-394-II-399 vol.2.
- [19] R. B. Rusu, Z. C. Marton, N. Blodow, and M. Beetz, "Persistent Point Feature Histograms for 3D Point Clouds," in Proc. of Int. Conf. on Intel Autonomous Systems, 2008, pp. 119-128.
- [20] S. Hickson, S. Birchfield, I. Essa, and H. Christensen, "Efficient Hierarchical Graph-Based Segmentation of RGBD Videos," in Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR), June 2014, pp. 344-351.
- [21] M. Grundmann, V. Kwatra, M. Han, and I. Essa, "Efficient Hierarchical Graph-Based Video Segmentation," in Proc. of IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR), June 2010, pp. 2141-2148.
- [22] A. L. Gibbs and F. E. Su, "On Choosing and Bounding Probability Metrics," International Statistical Review, vol. 70, no. 3, pp. 419-435, 2002.
- [23] L. Pardo, Statistical Inference Based on Divergence Measures. New York: Chapman & Hall/CRC, 2006, pp. 45–51.
- D. Holz, S. Holzer, R. B. Rusu, and S. Behnke, "Real-Time Plane Segmentation using RGB-D Cameras," in *Proc. of the 15th RoboCup* [24] International Symposium, vol. 7416, July 2011, pp. 307-317.
- [25] T. Lee, S. Lim, S. Lee, S. An, and S. young Oh, "Indoor Mapping Using Planes Extracted from Noisy RGB-D Sensors," in Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Oct 2012, pp. 1727-1733.
- [26] A. Uckermann, R. Haschke, and H. Ritter, "Realtime 3D Segmentation for Human-Robot Interaction," in Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Nov 2013, pp. 2136-2143.
- C. Tomasi and R. Manduchi, "Bilateral Filtering for Gray and Color [27] Images," in Proc. of Int. Conf. on Computer Vision (ICCV), Jan 1998, pp. 839-846.
- [28] C. E. Leiserson, R. L. Rivest, C. Stein, and T. H. Cormen, Introduction to Algorithms. MIT Press, 2001, ch. 21.
- J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A [29] Benchmark for the Evaluation of RGB-D SLAM Systems," in Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Oct 2012, pp. 573-580.
- [30] A. Handa, T. Whelan, J. McDonald, and A. Davison, "A Benchmark for RGB-D Visual Odometry, 3D Reconstruction and SLAM," in Proc. of IEEE Int. Conf. on Robotics and Automation (ICRA), May 2014, pp. 1524-1531.
## Prototyping during the requirements elicitation process in the development of an underground unmanned aerial system.

Jeremy Green, A Marnewick and JHC Pretorius<sup>1</sup>

Abstract—Prototyping of subsystem and system components is most often thought of as a development task. This paper shows the usefulness of prototyping as an activity in the requirements elicitation process, prior to any developement activities. It is approached from the field of engineering and technology management. It uses the Requirements Engineering approach to identify tools and methods for the development of the requirements for an underground unmanned aerial system for use in South Africa's' gold mines to inspect box-holes and ore-passes. Box-holes and ore-passes are vertical tunnels through which the ore must pass in moving from the stope, where it is mined, to the shaft, where it is hauled to the surface for processing. The more familiar new product development framework is compared to the requirement engineering process. The prototypes of a number of subsystems are presented, namely, a quadrotor platform, a platform preservation sensor array, an optical flow sensor for position holding, a vision sensor for operator visualization, and an operator interface. The perceived significant technological challenges are discussed as motivation in the choice of these subsystem prototypes that will be used in the interviews that are to form the basis of the requirements elicitation activity.

#### I. INTRODUCTION

In the research field of engineering and technology management, prototyping occurs at two phases. The more common use is as a technology demonstrator during the project development phase. Typically there will be many prototypes during this phase. However, this paper focusses on the other use of prototypes, that of during the *requirements elicitation* phase. At this stage there is no development team yet. The problem is still being understood and the requirements are being discovered. Thus enabling the formulation of the requirements specification that will be used to create and direct the development team [1]. The prototypes discussed in this paper are used not to demonstrate technologies (or solutions), but to encourage discussion about, and gain insight into, the problem, and improve the general understanding. Thus in this mining case study, there is no testing of the systems in a mine environment yet.

Section II discusses the background from three perspectives, mining, requirements elicitation, and requirement classification. Section III discusses the case study, specifically the requirements engineering tasks, technical challenges and the proposed subsystem prototypes to be used in the requirements elicitation process. Section IV then expands on each



Fig. 1. Gold Mine Structure

of the subsystem prototypes. Section V concludes the paper with conclusions and proposed future work.

#### II. PROJECT BACKGROUND

#### A. Mining

The project under discussion in this paper is of the development of an Unmanned Aerial/Aircraft System (UAS) (a quadrotor) for use in the inspecting box-holes and ore-passes in underground gold mining in South Africa. The project developed from a mine rescue conference workshop session about how robots could assist in mine rescue situations [2]. Because the acquisition of emergency equipment is hard to justify financially, an application was found that would also have benefits in a routine production environment. The basic functional requirements were documented for a case where a machine can inspect the vertical voids during production, as well as when they become periodically blocked creating an the emergency situation [3]. Thus releasing people from such dangerous jobs that have in the past resulted in fatalities [4].

1) Structure of a Gold Mine: The structure of a gold deposit and mine is shown in Figure 1. The gold ore deposit is called a reef. It is a narrow vein of ore ranging from several centimeters to a couple of meters thick. The reef dips from surface at between  $18^{\circ}$  and  $25^{\circ}$ , and plunges to unknown depths into the earth, while being 100's of kilometers in breadth. Current gold mines are considered

<sup>\*</sup>This work was not supported by any organization

<sup>&</sup>lt;sup>1</sup>Postgraduate School of Engineering Management in the Faculty of Engineering and the Built Environment at the University of Johannesburg jjgreen@outlook.com, amarnewick@uj.ac.za, jhcpretorius@uj.ac.za



Fig. 2. Diagram of a box-hole structure and dimensions.

very deep, ranging from 1.5 km to 4 km underground. They are getting more dangerous to work, as well as more difficult to work, with the increased temperatures and rock stresses that accompany such deep workings, as well as additional costs of hauling the low yield ore further to the surface for processing.

A vertical shaft, or elevator, is used to access the mining depth with access levels approximately 200 m apart. Horizontal traveling ways (also called haulages or access tunnels) are used to access the deposit from the shaft, and can be tens of kilometers long, accessed by means of a railway system. An ore-pass (not shown in the Figure) runs parallel to the shaft. Ore from each level is deposited into the ore-pass, where it falls down to the lowest level of the mine for loading into the cage and transported to surface for processing.

At the reef intersection, the mine structure changes to match the dip of the reef. *Raise* tunnels are developed along the reef plane, and horizontal tunnels called gullies (not shown) enable access to the stope where mining occurs. Ore is scraped by hydraulic scrapers down the stope, along the gullies and then down the raise to the level intersection, where a box-hole is used to load the rail car.

Figure 2 shows the typical configuration of a box-hole. The box-hole links the raise and the haulage tunnel, and enables ore to be loaded into the rail cars (hopper) for transport from the reef to the shaft, where it is dumped into the ore-pass. A typical box-hole is approximately 35 m long and 2 x 2 m square in section. An upper vertical section is capped with a screen (grizzley) with a 30 cm x 30 cm aperture, to prevent large rocks from entering and blocking the chute. The lower section is at 50° to reduce the kinetic energy of the falling ore, and is capped by a 'box-end' which controls the flow of ore for the loading of the hopper.



Fig. 3. Requirements process [9].

#### B. Requirements Elicitation

Getting the requirements right is a fundamental step in ensuring a successful research project execution. Multiple input sources are interrogated to understand the need, which is then analyzed, documented and verified with the stakeholders to create an agreed set of deliverables, the System Requirements Specification (SyRS), [5] [6]. In New Product Development (NPD) and system development, the Requirements Engineering (RE) process is the same as it is for software engineering and business analysis, as in Figure 3.

The four steps of:

- 1) elicitation
- 2) analysis
- 3) documentation
- 4) verification

are common across disciplines, however, the techniques employed vary amongst the project types. There are many books written about the subject [7]. In [8] the NPD process is described as in Figure 4. The requirements engineering process maps to the concept development phase, combining the steps of 'identify customer needs' through to 'set final specifications'.

Requirements elicitation is the process of gaining an understanding of the customers and users needs for the planned system and their expectations of it [10]. [10] goes on to define prototyping as a quick and rough (i.e. incomplete, untested and potentially flawed) version of a subsystem. Its purpose is to provide a physical artifact, around which discussions can occur that lead to a better understanding of the subsystem, and its required capability, by all stakeholders, both customers and future developers. [11] motivates for using rapid prototyping during the elicitation stage as an effective tool for acquiring information and knowledge about a new system or product (as opposed to analyzing and understanding an existing system or problem).

## C. Minimum Viable Technology vs. Commercially Viable System

In the discussion during the requirements elicitation process, it is typically the final system that is discussed and



Fig. 4. New Product Development and Requirements Engineering combined

envisaged. In this text it is referred to as the Commercially Viable System (CVS). However, in a technology development project, there are interim development steps that are executed during the project. These project phases, or stages, will generate different prototypes. It is important to note that this is different to the prototypes discussed in relation to the requirements elicitation phase in this paper. One way of grading these prototypes is a Technology Readiness Level (TRL) progression. Initially developed by NASA for the development of complex systems like spacecraft [12], [13] defines TRL as follows: "it is a discipline-independent, programmatic figure of merit (FOM) to allow more effective assessment of, and communication regarding the maturity of new technologies". TRL's have gained much support and have been adopted by the United States Department of Defense [14], and other large research organizations [15]. The CVS may map to a TRL 8 or 9.

In this text we refer to the first system that is to be developed as the Minimum Viable Technology (MVT). These terms are discussed further in [16], but broadly speaking, the MVT represents a degraded subset of requirements from the CVS system. The MVT demonstrates a capability, reduces or mitigates some technical risk, clarifies the problem and solution by presenting a possible system, thereby enabling a better understanding of the CVS requirements. The MVT may map to a TRL 4, 5 or 6.

#### III. REQUIREMENTS ELICITATION FOR THIS CASE STUDY

The requirements elicitation methods chosen for this project, as well as the justification for their choice are discussed in [16]. Interviews are the primary method, with

the use of a questionnaire, brain storming, scenarios, use cases, and prototypes to prompt discussion to discover the requirements for both CVS and MVT. A domain specialist is used to create a baseline set of initial requirements. These initial requirements will be presented in the context of user scenarios for discussion during the interviews.

#### A. Technological Challenges

There are a number of significant technological challenges that were identified in the background discussed in section II. Briefly, they are:

- **Platform preservation system:** To stop the aerial platform from flying into a wall, floor or ceiling, even if the operator inadvertently tries to fly it in that direction.
- User interface: Identifying the actual end user is to be a significant outcome of the elicitation process. It could be an unskilled mining operator, or it could be a specialized pilot, depending on the deployment model chosen. If the mine were to be the owner and operator of the hardware, it will be a task delegated downwards, potentially to an unskilled miner. If however, the system was deployed as a service by a specialised company, the mine would pay for the data resulting from the 'flight', and it will likely be a skilled operator. These two scenarios could well result in different requirements for the graphical user interface (GUI). In either case however, it is important to determine what information the operator needs/wants, and how the operator would like to transfer instructions to the aerial platform, i.e. control the Unmanned Aircraft (UA).



Fig. 5. Small, simple cheap quadrotor demonstration platform

- Determining safe flight zone: The operator will not always have visual line-of-sight (VLOS) to the platform, and will need to teleoperate the vehicle using only the GUI. There will therefore need to be some data processing to assist the operator to determine where a safe zone for flight is.
- Drift due to ventilation air flow: The aerial platform could drift down the passage due to the ventilation air without significant change in the sensor data or in the GUI. Unplanned movement is undesirable, as the platform should only move based upon an operator instruction.

Based upon these identified challenges, prototypes have been developed/proposed for use in the elicitation process.

#### B. Subsystem Prototypes

Subsystem prototypes are to be used to generate discussion about technical issues and possibilities for addressing the challenges. The prototypes do not represent the actual solutions, but rather are used to indicate some technology capability as well as to generate discussion about the required CVS capabilities. The subsystem prototypes are:

- A basic quadrotor platform.
- An ultrasonic array as a platform preservation system.
- An optical flow sensor as a possible way to overcome drift.
- An ASUS Xtion Pro live for visualization and depth analysis to determine the access potential for the platform. i.e. will it fit?
- GUI, an illustration of how the operator could interact with and control the platform.

The prototypes are discussed further in the following section.

#### IV. PROTOTYPES

The five subsystem prototypes are now discussed in more detail .



Fig. 6. Ultrasonic obstacle detection array

#### A. Platform

The use of a quadrotor platform appears obvious in this instance, but some discussion is perhaps warranted. The possible platforms are ground, suspended and aerial [17]. As the intended application is in a near vertical environment, or in a cluttered floor environment, the use of a ground vehicle is unsuitable. The use of a suspended platform is feasible for top entry to the ore-pass and box-hole. However, in the case of a blocked chute, it is necessary to gain access from below to determine the position of the blockage. Access from below can only be achieved with the use of an aerial vehicle with hovering and vertical take off an landing (VTOL) capabilities. Thus a small, simple and cheap quadrotor has been chosen as a discussion piece for the interviews, shown in Figure 5.

#### B. Platform Preservation System

The platform preservation system for obstacle detection/avoidance sub-system prototype that has been built (See Figure 6) is based upon an Arduino Uno and the HC-SR04 ultrasonic sensor [18]. An array of 10 sensors was initially intended, however, limitations in the arduino I/O has resulted in the initial prototype having six sensors that are sequentially polled with a 50  $\mu$ s timeout to avoid crosstalk. The ultrasonic sensor sends out a 40 MHz 'ping' and measures the time taken for the sound to return as a reflection off an object. It has a 15° field of view. The cycle time for polling the sensors is dependent upon the response time of each sensor, which is dependent on the distance measured which is dependent of the environment around the sensor system. The cycle time for polling the 6 sensors in a 4 m x 4 m room varied between 70  $\mu$ s and 90  $\mu$ s, implying a ten sensor system would be 116  $\mu$ s to 150  $\mu$ s.



Fig. 7. Asus Xtion Live sensor

It must be noted that the intended platform preservation system will be a three-dimensional system. Upward and downward facing sensors will prevent collision with the ceiling (hanging wall) and floor (foot-wall), and/or maintain a constant position between the hanging and foot walls. While this prototype is a coplanar system designed to preserve the platform from collisions when moving left, right, forward and backwards.

#### C. Operator Visualisation Sensor

For the operator to 'see' where the platform might move to, a Red, Green, Blue, Depth (rgbd) sensor has been chosen for the prototype. The Asus Xtion Pro Live [19] is an open source sensor that has been used in the past for similar visualisation activities in underground gold mines [20]. Figure 7 shows a disassembled Asus Xtion Pro Live, a 480x360 resolution rgbd sensor. The prototype uses Open NI and has a 0.8 m to 3.5 m range at 30 frames per second (fps), sufficient for the mine tunnel environment.

Figure 8 shows a depth map that can be used to determine where the platform can safely fly in a tunnel environment. The intention is to limit the operator instructions to those areas/directions that are safe. With a forward facing sensor, the platform will only be able to progress in the direction that the sensor is facing. 'Forward' will be different for different



Fig. 8. Tunnel depth map from Xtion sensor from [21]



Fig. 9. optical flow sensor from 3D Robotics for \$150

deployment scenarios. For example, in a tunnel, the sensor will point horizontal; in an ore-pass, the sensor will point vertically up or down; in an intermediate slope (raise or stope), the sensor tilts to match the proposed direction of travel for the platform, either upslope or down slope. There is no "backwards". The platform must rotate, tilt the sensor, determine if it will fit (through image analysis), then fly 'forward' in the direction that the sensor is pointing.

#### D. Drift Sensor

Figure 9 shows an optical flow sensor from 3D Robotics [22]. The PX4FLOW (Optical Flow) Sensor is a specialized high resolution downward pointing camera module that uses the ground texture and visible features and a rangefinder to determine aircraft ground velocity. [23] has shown the potential for combatting drift with such a sensor. [24] provides a survey of techniques and hardware that can be employed. It indicates that while none have used it specifically for position hold implementations, it has been effective on VTOL platforms for obstacle avoidance, terrain following, vertical landing, velocity estimation, and visual odometry. Some additional work would be needed to develop a prototype specifically for this application, to combat drift due to crosswinds from the ventilation air flowing down the tunnels. No system is proposed for this prototype, just a discussion about the sensor capabilities and the problem requirements. This discussion will enable the discovery of the system requirements for the CVS and MVT.

#### E. Operator Interface

The operator interface GUI will be on a portable computer. At lease some of the flight will be executed without VLOS of the aerial platform. Therefore, there needs to be sufficient information on the GUI for the operator to be able to make decisions about what to do. Figure 10 is a sketch of one such possibility. Using sketches is a simplistic first step in engaging potentially non-computer literate stakeholders, like miners, without intimidating them. Thus enabling them to easily add their thoughts, and enabling the capture of their inputs into what is, and is not, needed in the GUI. Proposing a GUI prototype will prompt discussion about



Fig. 10. Simple illustration of possible operator interface

a number of items: who the operator will be; what the operator environment will be like; how the operator will make decisions; what data/information they would need to make those decisions; how that information is to be displayed or conveyed to the operator such that it is unambiguous and useful. The logical next step is to develop the prototype on a computer system for the stake holders to interact with, and provide feedback on.

Typical GUI would include the readings from the ultra sonic sensor array displayed as a modified bar chart. Also, the sensor depth data can be analyzed to determine if the platform is dangerously close to an obstacle or wall. The display then colored to indicate the obstacle proximity (see Figure 10). Another example is that the rgbd sensor data are analyzed to indicate the possible trajectories that the platform can take. A green frame around the image indicates a feasible forward trajectory, a red frame indicates a blocked forward path, and the necessity to change the platform orientation and/or position, by either a left/right rotation or up/down movement, to find a clear forward path.

#### V. CONCLUSIONS

In this paper we have discussed the requirements elicitation prototypes to be used in the development of an UAS for use in South Africa's underground gold mines for inspecting ore-passes and box-holes. A summary of gold mining is given, explaining the challenges, and a background to the project is presented, outlining how this application was chosen for investigation. A discussion on requirements engineering and new product development processes precedes motivation for how prototyping can be a valuable tool in the elicitation process. The significant technical challenges for a UAS in an underground mining environment were outlined. Five subsystem prototypes were described that would be used in the requirements elicitation process for the underground UAS for box-hole and ore-pass inspection. The prototypes will be used in the interview discussions to assist in determining what exactly a solution system needs to achieve, as well as to more fully understand the deployment environment, and how that environment will effect the solution.

Follow up work includes the completion of an accompanying questionnaire to lead the interviews, and enable comparison results from a variety of stakeholder interviews. The stakeholder network will classify the requirements into MVT and CVS requirements, and this classification is to be mapped onto the TRL framework.

#### REFERENCES

- A. Sutcliffe and P. Sawyer, "Requirements elicitation: Towards the unknown unknowns," in *Requirements Engineering Conference (RE)*, 2013 21st IEEE International, pp. 92–104, July 2013.
- [2] J. Green, "Mine rescue robot, every mine needs a robot?," in 3rd Annual Mine site Emergency Preparedness and Rescue Innovation, 2013.
- [3] J. Green, "Mine rescue robots requirements outcomes from an industry workshop," in *Robotics and Mechatronics Conference (RobMech)*, 2013 6th, pp. 111–116, IEEE, 2013.
- [4] T. Stacey and B. Erasmus, "Setting the scene rockpass accident statistics and general guidelines for the design of rockpasses," *Journal South African Institute of Mining and Metallurgy*, vol. 105, no. 11, p. 745, 2005.
- [5] A. Marnewick, A Socio-Economic View of the Requirements Engineering Process. PhD thesis, University of Johannesburg, 2013.
- [6] A. Marnewick, J. Pretorius, and L. Pretorius, "A south african perspective of the requirements discipline: An industry review," SAIEE Africa Research Journal, vol. 105, no. 3, pp. 112–126, 2014.
- [7] D. Zowghi and C. Coulin, "Requirements elicitation: A survey of techniques, approaches, and tools," in *Engineering and Managing Software Requirements* (A. Aurum and C. Wohlin, eds.), ch. 2, pp. 19– 46, Springer Berlin Heidelberg, 2005.
- [8] K. T. Ulrich and S. D. Eppinger, Product design and development. New York: McGraw-Hill, 2003.
- [9] G. Kotonya and I. Sommerville, *Requirements Engineering: Processes and Techniques*. Chichester, UK: John Wiley and Sons, 1998.
- [10] R. R. Young, *The requirements engineering handbook*. Artech House, 2004.
- [11] N. Maiden and G. Rugg, "Acre: Selecting methods for requirements acquisition," *Software Engineering Journal*, vol. 11, no. 3, pp. 183– 192, 1996.
- [12] J. C. Mankins, "Technology readiness levels," White Paper, April, vol. 6, 1995.
- [13] J. C. Mankins, "Technology readiness assessments: A retrospective," Acta Astronautica, vol. 65, no. 9–10, pp. 1216 – 1223, 2009.
- [14] U. DoD, "Technology readiness assessment (tra) guidance," *Revision posted*, vol. 13, 2011.
- [15] D. W. Engel, A. C. Dalton, K. K. Anderson, C. Sivaramakrishnan, and C. Lansing, *Development of technology readiness level (TRL) metrics* and risk Measures. Pacific Northwest National Laboratory, 2012.
- [16] J. Green, A. Marnewick, and J.-H. Pretorius, "Requirements degradation for the creation of a first prototype," in *Management of Engineering and Technology (PICMET), Portland International Conference on*, PICMET, August 2015.
- [17] R. R. Murphy, J. Kravitz, S. L. Stover, and R. Shoureshi, "Mobile robots in mine rescue and recovery," *Robotics & Automation Magazine, IEEE*, vol. 16, no. 2, pp. 91–103, 2009.
- [18] Cytron, "http://www.cytron.com.my/p-sn-hc-sr04."
- [19] "https://www.asus.com/us/multimedia/xtionprolive/specifications/," 27 August 2015.
- [20] J. Dickens and M. Price, "The design of an automated 3d-thermal mine scanning tool," in *Robotics and Mechatronics Conference of South Africa (ROBOMECH), 2012 5th*, pp. 1–6, Nov 2012.
- [21] J. Green, "Small machines in deep mines, sensors for autonomous machines in underground mining," in *National Instruments Technical Symposium*, March 2012.
- [22] 3dRobotics, "https://store.3drobotics.com/products/px4flow," August 31 2015.
- [23] J. Kim and G. Brambley, "Dual optic-flow integrated navigation for small-scale flying robots," *Proc. of Australasian Conference on Robotics and Automation, Brisbane, Australia*, 2007.
- [24] H. Chao, Y. Gu, and M. Napolitano, "A survey of optical flow techniques for uav navigation applications," in *Unmanned Aircraft Systems (ICUAS)*, 2013 International Conference on, pp. 710–716, May 2013.

## Single-labelled Music Genre Classification Using Content-Based Features

Ritesh Ajoodha, Richard Klein, and Benjamin Rosman

*Abstract*—In this paper we use content-based features to perform automatic classification of music pieces into genres.

We categorise these features into four groups: features extracted from the Fourier transform's magnitude spectrum, features designed to inform on tempo, pitch-related features, and chordal features.

We perform a novel and thorough exploration of classification performance for different feature representations, including the mean and standard deviation of its distribution, by a histogram of various bin sizes, and using mel-frequency cepstral coefficients.

Finally, the paper uses information gain ranking to present a pruned feature vector used by six off-the-shelf classifiers. Logistic regression achieves the best performance with an 81% accuracy on 10 GTZAN genres.

*Index Terms*—Music genre classification, feature selection, feature representation, MFCC aggregation, area moments, tempo detection, pitch detection, chordal identification, information gain ranking.

#### I. INTRODUCTION

**M**USIC genre, while often being vaguely specified, is perhaps the most common classification scheme used to distinguish music. Although single human responses to genre classification can be biased and stereotypical, there exists a consensus of broad genre definitions across populations worldwide.

Genre classification is one of multiple music classification methods, including mood and artist classification. Although these methods are also similarity-based measures across different music meta-data (e.g. lyrics, artist, timbre), genre offers a culturally authorised prominence on the construction of traditional classes which is more functional for music classification.

Music genre has such a pressing influence on consumers that a listener may prefer one song to another based more on the song's genre than the actual song itself [1] [2]. End-users are more likely to browse music by genre than artist similarity, recommendation, or even music similarity [3]. Therefore, successful music genre classification algorithms will enable users to browse music within genre categories.

Our aim is to explore the space of automatic music genre classification, so as to decrease search-time for music pieces

R. Ajoodha is with the Department of Computer Science and Applied Mathematics, The University of the Witwatersrand, Johannesburg, South Africa. e-mail: ritesh.ajoodha@gmail.com.

R. Klein is with the Department of Computer Science and Applied Mathematics, The University of the Witwatersrand, Johannesburg, South Africa. e-mail: richard.klein@wits.ac.za.

B. Rosman is with the Council for Scientific and Industrial Research as well as the Department of Computer Science and Applied Mathematics, The University of the Witwatersrand, South Africa. e-mail: brosman@csir.co.za.

within large databases. This system relies only on the audio signal itself and does not consider any meta-data.

Observing interactions between genre classes through content-based features can unveil cultural associations that exist between these genre classes and is of musicological significance [4].

#### II. BACKGROUND

*Music Genre Classification* is the process of categorising music pieces using traditional and cultural aspects. These traditions and cultures are not precisely defined and so over the years it has become vague as to what characteristics secure music to a particular genre.

Traditional musical aspects are given by four characteristics [5]: melody, harmony, rhythm and sound (timbre, dynamics, and texture) which are hypothesised to contribute considerably to the notion of musical genre. However, standard genre textbook definitions are qualitative, subjective, context dependent, and therefore are difficult to automate.

As a result of the ambiguities that exist between contentbased genre definitions the ground truth classification accuracy becomes inescapably bounded as many people may disagree on a particular genre classification of a piece of music.

Composers often do not abide by "genre definitions", which makes us question whether some composers are accepted by currently "defined" music genres. For this reason music genre classification is categorised using human discretion and is therefore prone to errors and subjectivity as many pieces of music sit on boundaries between genres [6].

Successful genre classification makes use of cultural-based rather than content-based feature dissimilarity between genre classes. In this work we do not consider such meta-data, due to lack of availability.

Many genres do not only sound similar but also contain multiple sub-genres which share some similar characteristics. The difficulty of genre classification thereby increases when considering hundreds of other genre types and their respective sub-genres.

Although some authors provide an awareness of genre classification performance bounds imposed by human responses to genre classification [7] [8], further study in experimental research is needed to draw more concise conclusions regarding human responses to genre classification and how this affects ground truth. Humans are biased and subjective in genre classification, which ultimately leads to a lack of consensus in genre labels and thus poor quality of ground truth.

978-1-4673-7450-7/15/\$31.00 ©2015 IEEE

To make matters worse, genre definitions evolve over time and continuously give rise to newer structures that have significantly different feature compositions [9]. Therefore, regardless of feature dimensionality, well-built classification procedures are required to classify features successfully with some regard for genre development. This is difficult as the scalability of even the most powerful supervised classification models are unsatisfactory [10].

Benetos and Kotropoulos (2008)	75.0%
Bergstra et al. (2006)	82.5%
Holzapfel and Stylianou (2008)	74.0%
Li et al. (2003)	79.7%
Lidy et al. (2007)	76.8%
Panagakis et al. (2008)	78.2%
Sturm (2013)	83.0%
Tzanetakis and Cook (2002)	61.0%

TABLE I: Noteworthy genre classification algorithms on 10 GTZAN genres.

A review of the literature shows very few capable genre classification systems using the GTZAN dataset. The GTZAN dataset is a collection of 1000 thirty second excepts. The 1000 music excepts are categorised into 10 genres, 100 excepts for each genre. Systems thus far have not adopted automatic genre classification models for media retrieval and recommendation. Successful genre classification includes work by Sturm (2013) [11] who achieved 73-83% on 10 genres; and Bergstra et. al. (2006) [12] who achieved 82.50% on 10 GTZAN genres [13]. Table I shows some noteworthy music genre classification algorithms on 10 GTZAN genres.

#### **III. FEATURE ANALYSIS**

Music comprises of instrument sounds, speech sound, and environmental sounds [14] [15]. In this section we present several features that are hypothesised to be characteristics that can be used to correctly classify musical genre.

These features are organised into four main categories: *Magnitude-based features*, where timbral features that describe loudness, noisiness, compactness, etc. are presented; *Tempo-based features*, where methods that explore rhythmic aspects of the signal are provided; *Pitch-based features*, where algorithms that describe the pitch of music signals are presented; and finally, *Chordal Progression features*, where we explore chroma as a chordal (environmental) distinguishing feature.

Before we present these four families of features, we will firstly introduce four feature representations that will be explored for each feature distribution.

#### A. Feature Representation

In addition to the **mean**, the following feature representations will be applied to each feature and the best representation for each feature will be used in the final classification.

**The Feature Histogram:** The feature histogram arranges the feature's local window intensities into bin ranges. The content of each bin is counted and modelled by a frequency histogram. The histogram bin values are normalised and used for classification. **MFCC Aggregation:** MFCC representation is a wellknown feature representation that takes the first n MFC coefficients (coefficients that make up the short-term power spectrum of sound) as it would a 16khz signal [16] [17]. If the feature contains more than one dimension, then each dimension is assessed independently and n coefficients will be produced per dimension. In this paper we set n = 4.

Area Moments: Image moments is a central concept in computer vision and has its root in image processing. Fujinaga (1996) [16] produced 10 such moments for image processing: an image is treated as a 2-dimensional function f(x, y) = z, where x and y are indexes of the underlying matrix. The feature values extracted from the audio signal will be treated as a 2-dimensional image and Fujinaga's moments algorithm will be applied to the feature vector.

#### B. Magnitude-based Features

The magnitude spectrum, obtained from the fast Fourier transform of a signal, houses a family of spectral features which can be used for genre classification. Exploration of the magnitude spectrum has allowed us to identify signal change, noisiness, loudness and many other spectral features that describe aspects of discrete time signals for automatic music genre classification.

Exploring peak-based features, from the local maxima of the frequency domain, creates opportunities to analyse the signal more thoroughly. In this section we explore magnitude-based features for music genre classification.

1) Spectral Slope: The spectral slope can be observed when natural audio signals tend to have less energy at high frequencies. Peeters (2004) [18] provides a way to quantify this by applying a linear regression to the *magnitude spectrum* of the signal, which produces a single number indicating the slope of the line-of-best-fit through the spectral data.

2) *Compactness:* Compactness is a measure of the noisiness of a signal [17] and is calculated by comparing the value of a magnitude spectrum bin with its surrounding values. In many genres (e.g. metal) a random and persistent disturbance that obscures the clarity of sound is desired, which this feature will detect. Figure 3(a) shows the compactness feature values distributed over 10 GTZAN genres.

*3) Loudness:* Specific loudness is the loudness associated with each critical band of hearing. Total loudness has been used for multi-speaker speech activity detection, automatic speech recognition, instrument recognition and music genre classification.

4) Onset Detection: Onset detection describes information about the initial magnitude of a piece of music [19]. This feature describes the rise in magnitude from zero to some initial value.

5) *Peak Detection:* Studying the peaks of a signal allows us to account for various principal features that are contained within a signal. For example, peak-based features such as crest factor, peak flux, centroid, and smoothness can help us describe the quality of AC waveform power and detecting vibration. The peak detection algorithm by [20] will be used for extracting peak-based features. Mckay (2005) calculated peaks by detecting local maxima in the frequency bins, and these maximum are calculated within a threshold where the largest maxima within this threshold is considered [20]. These global peaks per threshold are considered without any information about bin location. In our experiments we took a peak threshold of 10. Treating this set of peak values together as a 16khz signal, we then represent these peak values by the centroid, flux, and smoothness features.

6) Spectral Flux: Spectral flux is a content-based feature that measures the rate of change of the magnitude spectrum. This is achieved by comparing every frame of the magnitude spectrum with its previous frame.

7) Spectral Variability: Statistical variability measures dispersion in data, i.e. how closely or spread-out the signal is clustered. We can achieve this by measuring the standard deviation of the magnitude spectrum of the signal.

8) *Mel-Frequency Cepstral Coefficients:* Mel-frequency cepstral coefficients (MFCCs) are the coefficients that together make up a Mel-frequency cepstrum. The components of MFCC are those from the cepstral representation of the audio signal. In the Mel-frequency cepstrum the frequency bands are equally spaced which favours the human auditory system more than using the cepstrum feature alone, which uses linearly-spaced frequency bands.

*9) Spectral Flatness:* Spectral flatness is a feature used to calibrate how pure tonal sounds are in comparison to noisy ones. Pure tonal sound refers to resonant structure in a power spectrum, compared to other parts containing white noise.

10) Spectral Rolloff: According to [18], [12], spectral rolloff point is the frequency such that 85% of the signal energy is contained below this frequency. It is correlated with the harmonic/noise cutting frequency [18]. Figure 3(d) shows the spectral rolloff feature values distributed over 10 GTZAN genres.

#### C. Tempo Detection

Most music retains regular rhythmic formations that creates an impression of tempo. With the purpose of understanding the nature of music to perform genre classification, tempo must be understood and preserved as a feature description. In this section we establish tempo detection schemes for music genre classification. Having already established a method to detect the vitality in a music excerpt by using spectral energy, which is the root mean square (RMS) of the music signal, we present in this section the beat histogram as a crucial feature vector.

1) Energy: Energy is a fundamental descriptor used in speech and audio processing [21]. Energy is measured by calculating the RMS of a discrete-time signal. Figure 3(c) shows the energy feature values distributed over 10 GTZAN genres.

Examining the arithmetic average of the first n windows of a signal (for our experiments we took n = 100) and calculating the fraction of these which are below the average, we can calculate the percentage of silence that exists in the signal - as the *fraction of low energy*.

2) The Beat Histogram: The beat histogram is an arrangement of signal strength to yield rhythmic intervals. This

is accomplished by measuring the energy of n consecutive windows and computing the fast Fourier transform of the result. This type of feature will produce a very large design matrix and so a simple feature representation is needed. In our experiments the mean feature representation outperformed MFCC and the 20-bin feature histogram.

#### D. Pitch and Speech Detection

Pitch is a perceived characteristic contained in the frequency of music content. Most music of the same genre exhibit melodies that are just combined notes from a scale set. For example, most notes from an impressionistic piece are taken from whole-tone scales, whereas notes from a jazz pieces of music are taken from pentatonic scales. However, often environmental sounds overtone pitch, disguising available pitch-related elements, which make it difficult to extract pitch computationally. Even human auditory systems can find it difficult to distinguish pitch under these conditions.

In this section we explore pitch and speech related algorithms as an amalgam of these characteristics are hypothesised to describe singing. Together, pitch and speech detection schemes can help us understand gliding, portamento, or even vibrato.

1) Amplitude Modulation: For many musical instruments amplitude periodic modulation is a distinctive quantity. Style introduces characteristic amplitude variation into music tones. It has been observed that changing amplitude envelopes leads to similarity judgments on musical timbre [22]. The energy envelope is useful to extract features measuring amplitude modulation (AM). It has been observed that heuristic strength and frequency of AM can be calculated at two frequency ranges: the first range is between 4 and 8 Hz (where the AM is in conjunction with vibrato) and the second range is between 10 to 40 Hz which correspond to "graininess" or "roughness" of the tone.

2) Zero Crossing Rate: The zero crossing rate (ZCR) is the frequency of sign changes that occur along a discrete-time signal. Being a thorough percussive descriptor, this feature has been used in both speech recognition as well as in audio information retrieval. Figure 3(b) shows the strongest beat values distributed over 10 GTZAN genres.

#### E. Chordal Progressions

Introducing spectral feature extraction to genre detection problems created opportunities to exploit single characteristics of music. Chord structure and progressions have defining traits of music for many years.

1) Chroma: Chroma is defined as a 12 component design matrix where each dimension represents the intensity associated with a particular semitone, regardless of octave [23]. This section implements MFCC-based chroma by extracting MFCCs derived from a 12 component chroma design matrix. Since the components of chroma describe the distribution of semitones in a piece of music, it also informs us how notes are arranged and thus provides information about chordal harmonies. Therefore, modelling chroma indicates if a particular genre displays an attachment or relation to harmonic chordal progressions, as some genres do.

#### **IV. FEATURE SELECTION**

In the upper part of Table II, we present the features mantained after using the Information gain ranking algorithm. Information gain ranking is a filter method that evaluates the worth of a feature by measuring the information gain with respect to the class. The lower part of Table II lists the eliminated features.

The cut-off point was chosen by considering Figure 1, which shows the results of taking different numbers of features with the highest contributions and using them to classify 10 GTZAN genres. The red-line in Figure 1 shows the cutoffpoint taken at 459 features in its respective representation. Figure 1 suggests we could have chosen about 100 features and achieved between 70-75% classification accuracy with minimal performance loss, but we extended this for robustness reasons.

#### V. AUTOMATIC MUSIC GENRE CLASSIFICATION

In this section we use the selected features outlined in the upper portion of Table II to perform genre classification on 10 GTZAN genres. Table III yields the results of this experiment: the first column lists the classifiers used; the second column gives us the accuracy for each classifier to correctly identify 10 genres; finally, the third column lists the time to build each classification model. The implementation details of each of the algorithms are outlined in Table IV. Six of-the-shelf classifiers were used: Naïve Bayes; Support Vector Machines; Multilayer Perceptron; Linear Logistic Regression Models; K-Nearest Neighbours; and Random Forests.



Fig. 1: Classification accuracy vs the number of highest contributing features to classify 10 GTZAN genres.

It is seen that all of the classification algorithms outperform the Naïve Bayes method. The Support Vector machine, Multilayer Perceptron and Random Forests are aligned by their performance with the Multilayer Perception taking the most time to build. The Linear Logistic Regression Model provides the best classification score of 81%. However, with the exception of the Multilayer Perceptron, the Linear Logistic Regression Model takes the longest time to build. Figure 2 shows the confusion matrix for 10 GTZAN genres using Linear Logistic Regression Models with 10-fold cross validation. The particular cluster overlap between rock and country

Features Maintained	Rep.	Dim. 459
Spectral Flux	MFCC	4
Spectral Variability	MFCC	4
Compactness	Mean + SD	2
MFCCs	MFCC	52
Peak Centroid	Mean + SD	2
Peak Smoothness	SD	1
Complex Domain Onset Detection	Mean	1
Loudness + Sharpness and Spread	Mean	26
OBSI + Radio	Mean	17
Spectral Decrease	Mean	1
Spectral Flattness	Mean	20
Spectral Slope	Mean	1
Shape Statistic spread	Mean	1
Spectral Centroid	MFCC	4
Spectral Rolloff	SD	1
Spectral Crest	Mean	19
Spectral Variation	Mean	1
Autocorrelation Coefficients	Mean	49
Amplitude Modulation	Mean	8
Zero Crossing + SF	MFCC	8
Envelope Statistic Spread	Mean	1
LPC and LSF	Mean	12
RMS	Mean + SD	2
Fraction of Low Energy	Mean	1
Beat Histogram	SD	171
Strength of Strongest Beat	Mean	1
Temporal Statistic Spread	Mean	1
Chroma	MFCC	48
Features Eliminated	Rep.	Dim. 223
Peak Flux	20-bin FH	20
Peak Smoothness	Mean	1
Shape Statistic centroid, skewness	Mean	1
Shape Statistic Kurtosis	Mean	2
Strongest Frequency of Centroid	MFCC	4
Spectral Rolloff	Mean	1
Strongest Frequency of FFT	MFCC	4
Envelope Centroid, Skewness and Kurtosis	Mean	4
Beat Histogram	Mean	171
Strongest Beat	Mean + SD	2
Strength of Strongest Beat	SD	1
Fraction of Low Energy	SD	1
Beat Sum	MFCC	4
Relative Difference Function	MFCC	4
Temporal Statistic Centroid	Mean	1
Temporal Statistic Skewness	Mean	1
Temporal Statistic Kurtosis	Mean	1

TABLE II: The features maintained (upper portion) and the eliminated features (lower portion). Column two and three list the feature representation and feature dimension respectively.

music (and rock and disco) is observed. Although our results are in line with the best performing methods, and we have not exceeded them, we offer a valuable contribution in the form of feature analysis and representation for music genre classification.

#### VI. CONCLUSION AND RECOMMENDATIONS

Although recent classification accuracy suggests that the performance of learning models for genre classification have become bounded, there is no confirmation to date to suggest these bounds cannot be exceeded. Nonetheless, small changes to existing models are unlikely to produce significantly better classification scores. Therefore, more attention to how feature extraction and classification are performed, or perhaps completely new approaches, are crucial to greatly exceed these bounds.

Classifier	Accuracy	Time to build model
Naïve Bayes	53.2%	0.56 sec
Support vector machines	75.4%	3.82 sec
Multilayer perceptron	75.2%	27.48 sec
Linear logistic regression models	81.00%	25.25 sec
K-nearest neighbours	72.80%	0.01 sec
Random forests	75.7%	18.08 sec

TABLE III: Automatic genre classification using the thinned feature vector.

Classifier	Parameters used
Nu D	
Naive Bayes	Used a normal distribution for numeric at-
	tributes and supervised discretization
Support vector ma-	Kernal degree = $3$ ; tolarance of termination
chines	criteria = $0.001$ ; epsilon for the loss function =
	0.1; did not normalise; used polynomial kernal:
	$(gamma * u'v + coef0)^{degree}$ .
Multilayer perceptron	Number of hidden layers = Number of classes;
	learning rate = $0.3$ ; training time = 500 epochs;
	validation threshold = $20$ .
Linear logistic regres-	Maximum number of iterations for LogitBoost
sion models	= 500
K-nearest neighbours	Number of neighbours to use $= 1$ ; using the ab-
C C	solute error for cross validation; Linear search
	algorithm
Random forests	Number of trees used = $1000$

TABLE IV: Implementation details of each classification algorithm.

Erroneous genre labels are often caused by inexperienced respondents and not being exposed to enough of the recording [7], [24], [8]. The reliability of a learning model is purely measured by the quality of its ground truth and so extensive measures must be taken to ensure that the ground truth is well founded and motivated.

Since genre classification is usually performed by humans who observe cultural features (observations of arts and other manifestations of genre cognitively regarded collectively) more than content related features, we should not expect to achieve ground breaking results by classifying genre purely on content-based features. This is evident as the best genre classification algorithms using content-based features only achieve between 75-83% on 10 GTZAN genres.

Incorporating cultural features with structural ones in the feature domain could notably increase current classification rates [25]. Large scale musical structures are present in most music genre types. Understanding the form (cyclic, binary, rondo) of a piece of music can immediately designate a small set of potential genre categories to which the piece could belong. These overall structure-based feature descriptions can be preserved in learning models by using classifiers that exhibit memory<sup>1</sup>. Preserving memory in learning models have been mostly ignored and could hold the key to better understanding chordal progressions and complex melodic structures.

The musicality of a listener is not only required when constructing ground truth, but can also be used to satisfy a particular customer's genre preference. Further empirical research in human responses to genre classification can reveal if certain consumers with different musicality will appreciate



Fig. 2: The confusion matrix for 10 GTZAN genres using linear logistic regression models with 10-fold cross validation. The row and column labels represent genre labels where:  $G_1$  = Blues,  $G_2$  = Classical,  $G_3$  = Country,  $G_4$  = Disco,  $G_5$  = Hiphop,  $G_6$  = Jazz,  $G_7$  = Metal,  $G_8$  = Pop,  $G_9$  = Reggae, and  $G_{10}$  = Rock.

music differently. Empirical research should compare and contrast different classification scores for different kinds of customers in terms of age, culture, and musicality. This type of psychological research will enhance our understanding of the possibilities to increase the dependability of ground truth and will also allow us to personalise multiple learning models to cater for groups of individuals' needs rather than forcing a one fits all approach.

#### ACKNOWLEDGMENT

The first author gratefully acknowledges the support of the NRF Scarce Skills Doctoral Scholarship, and the Post-Graduate Merit Award Scholarship for Doctoral Studies.

#### REFERENCES

- A. C. North and D. J. Hargreaves, "Liking for musical styles," *Musicae Scientiae*, vol. 1, no. 1, pp. 109–128, 1997.
- [2] H. G. Tekman and N. Hortacsu, "Aspects of stylistic knowledge: What are different styles like and why do we listen to them?" *Psychology of Music*, vol. 30, no. 1, pp. 28–47, 2002.
- [3] J. H. Lee and J. S. Downie, "Survey of music information needs, uses, and seeking behaviours: Preliminary findings." in *ISMIR*, vol. 2004, 2004, p. 5th.
- [4] C. McKay and I. Fujinaga, "Automatic music classification and the importance of instrument identification," in *Proceedings of the Conference* on Interdisciplinary Musicology, 2005.
- [5] J. D. White, *The analysis of music*. Prentice-Hall Englewood Cliffs, NJ, 1976.
- [6] T. Li, M. Ogihara, and Q. Li, "A comparative study on contentbased music genre classification," in *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval.* ACM, 2003, pp. 282–289.
- [7] R. O. Gjerdingen and D. Perrott, "Scanning the dial: The rapid recognition of music genres," *Journal of New Music Research*, vol. 37, no. 2, pp. 93–100, 2008.

<sup>&</sup>lt;sup>1</sup>Much like how hidden Markov models or recurrent neural networks work.



types (e.g. Blues, Classical, Jazz, and Country).

(a) The compactness feature ranges for all 10 GTZAN genres. Com- (b) The ZCR feature ranges or all 10 GTZAN genres. Distinguishes pactness distinguishes some genres particularly well from other genre metal, disco, and hiphop particularly well from other genres types.



(c) The energy feature ranges for all 10 GTZAN genres. Distinguishes (d) The spectral rolloff frequency ranges for all 10 GTZAN genres. classical, pop, hiphop, metal, and jazz particularly well from other Distinguishes metal particularly well from other genres types. genres types.

Fig. 3: Different features distinguishing a variety of GTZAN genres.

- [8] S. Lippens, J.-P. Martens, and T. De Mulder, "A comparison of human and automatic musical genre classification," in Acoustics, Speech, and Signal Processing, 2004. Proceedings.(ICASSP'04). IEEE International Conference on, vol. 4. IEEE, 2004, pp. iv-233.
- [9] M. Grimaldi, P. Cunningham, and A. Kokaram, "An evaluation of alternative feature selection strategies and ensemble techniques for classifying music," in Proc. Workshop on Multimedia Discovery and Mining. Citeseer, 2003.
- [10] C. McKay and I. Fujinaga, "Musical genre classification: Is it worth pursuing and how can it be improved?" in ISMIR, 2006, pp. 101-106.
- [11] B. L. Sturm, "On music genre classification via compressive sampling," in Multimedia and Expo (ICME), 2013 IEEE International Conference on. IEEE, 2013, pp. 1-6.
- [12] J. Bergstra, N. Casagrande, D. Erhan, D. Eck, and B. Kégl, "Aggregate features and adaboost for music classification," Machine Learning, vol. 65, no. 2-3, pp. 473-484, 2006.
- [13] B. L. Sturm, "The gtzan dataset: Its contents, its faults, their affects on evaluation, and its future use," arXiv preprint arXiv:1306.1461, 2013.
- E. Wold, T. Blum, D. Keislar, and J. Wheaten, "Content-based classifi-[14] cation, search, and retrieval of audio," MultiMedia, IEEE, vol. 3, no. 3, pp. 27-36, 1996.
- [15] R. Ajoodha, R. Klein, and M. Jakovljevic, "Using statistical models and evolutionary algorithms in algorithmic music composition," in The Encyclopedia of Information Science and Technology, 3rd ed., K.-P. Mehdi, Ed. Hershey, Pennsylvania, United States: IGI Global, 2014.
- [16] I. Fujinaga, "Adaptive optical music recognition," Ph.D. dissertation, McGill University, 1996.

- [17] C. McKay, R. Fiebrink, D. McEnnis, B. Li, and I. Fujinaga, "Ace: A framework for optimizing music classification." in ISMIR, 2005, pp. 42-49
- [18] G. Peeters, "A large set of audio features for sound description (similarity and classification) in the cuidado project," 2004.
- [19] C. Duxbury, J. P. Bello, M. Davies, M. Sandler et al., "Complex domain onset detection for musical signals," in Proc. Digital Audio Effects Workshop (DAFx), no. 1, 2003, pp. 6-9.
- C. McKay, I. Fujinaga, and P. Depalle, "jaudio: A feature extraction [20] library," in Proceedings of the International Conference on Music Information Retrieval, 2005, pp. 600-3.
- [21] L. Lu, H.-J. Zhang, and H. Jiang, "Content analysis for audio classification and segmentation," Speech and Audio Processing, IEEE Transactions on, vol. 10, no. 7, pp. 504-516, 2002.
- [22] P. Iverson and C. L. Krumhansl, "Isolating the dynamic attributes of musical timbrea," The Journal of the Acoustical Society of America, vol. 94, no. 5, pp. 2595-2603, 1993.
- [23] D. P. Ellis, "Classifying music audio with timbral and chroma features," in ISMIR 2007: Proceedings of the 8th International Conference on Music Information Retrieval: September 23-27, 2007, Vienna, Austria. Austrian Computer Society, 2007, pp. 339–340. D. Perrot and R. Gjerdigen, "Scanning the dial: An exploration of factors
- [24] in the identification of musical style," in Proceedings of the 1999 Society for Music Perception and Cognition, 1999, p. 88.
- [25] C. McKay, "Automatic genre classification of midi recordings," Ph.D. dissertation, McGill University, 2004.

2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech) Port Elizabeth, South Africa, November 26-27, 2015

## Self-Learning of Inverse Kinematics for Feedforward Control of Intracardiac Robotic Ablation Catheters

Ge Bian, Michael Lipowicz, Grant H. Kruger University of Michigan Ann Arbor, MI, USA. ghkruger@med.umich.edu

Abstract— This paper investigates the self-learning of inverse kinematics for the feed-forward control of a robot to position intracardiac catheters. Cardiac ablation is routinely performed to treat Atrial Fibrillation, and requires a catheter be accurately positioned in the heart, by hand or by a robot, under feedback control. This is typically a slow process and methods to reduce procedure times are needed. To investigate our proposed method, a robotic system to manipulate a standard intracardiac catheter was constructed. To safely develop our proposed learning system, a comprehensive dataset was collected using a magnetic tracking system to measure the catheter tip positions versus robot actuator positions. Initially, the robot began with no model of its kinematics. A Genetic Algorithm was used to decide on the next actuator sequence that would reduce the uncertainty in a Feedforward Neural Network (FFNN) based inverse kinematic model. An automated iterative process was followed where the robot would perform virtual experiments, to grow its knowledge of its inverse kinematics. After 791 learning cycles the final analysis revealed that the complete inverse kinematic relationship has been explored with the given constraints. A validation dataset indicated the learned FFNN model was able to predict x, y and z positions of the catheter tip to within  $\pm 0.17$  mm,  $\pm 0.73$  mm and  $\pm 0.62$  mm, respectively. The robot successfully self-learned its inverse kinematic model using the proposed methodology. Future work is required to investigate the influence of disturbances on positioning accuracy.

#### Keywords—catheter ablation, robot, learning

#### I. INTRODUCTION

Atrial Fibrillation (AF) is the most common sustained cardiac arrhythmia, affecting over 2 million people, and is associated with increased risk of death, stroke, heart failure and reduced quality of life [1-3]. Over the last decade, studies have found increasing hospitalization rates and treatment costs [4, 5]. Robotics provides an opportunity to reduce the training and procedure time, costs and complications [6].

Fig. 1 shows a cross-sectional diagram of a human heart. A healthy heart beat originates from the sinoatrial (SA) node in the right atrium with an interval of approximately 1 Hz. The electrical impulse causes contraction of the right atrium walls. Similarly, the left atrium is excited, as well as the atrioventricular (AV) node. The AV node slows the electrical signal, to allow the ventricles to fill with blood as the atria contract. Thereafter, the ventricles contract causing blood to be pumped to the rest of the body. AF is caused by re-entrant electrical properties, which are excited by a trigger site. This

results in high frequency, disorganized, cyclic contractions of the atrium [1, 7]. Drug therapies are available to control the arrhythmia, but have shown limited efficacy (40 - 60 %) and can exhibit serious side effects. Catheter ablation procedures have shown increased efficacy, reduced hospitalizations, and reduced treatment costs compared to drug therapy [3, 5, 8].



Fig. 1. Diagram of heart indicating catheter positioning.

To perform a catheter ablation procedure, a long, thin, flexible catheter is inserted through a sheath into the femoral vein near the groin and moved toward the heart. Once inside the atrium the flexible catheter can be positioned against the wall to identify the source of the AF by measuring electrical activity, as shown in Fig. 1. Once the source has been located the ablation can be performed, inactivating the tissue and blocking the conduction pathway [4]. Ablation is typically performed by increasing the heat of the tissue using electrical, laser energy, or through cryogenic cooling [3].

Positioning of the catheters are still predominantly performed through manual control of the catheter by a skilled electrophysiologist, using fluoroscopic X-ray or computerized tomography (CT) guidance for position feedback (aided by intracardiac ultrasound) [5, 9]. The catheter tip can be positioned by moving the catheter axially in a sheath, bending the distal end using a pull-wire connected to a control on the handle, or rotated by rotating the handle. These procedures are therefore highly dependent on operator skill. In addition to the increased risk of non-ergonomic strain on the physicians due to the manual nature of the procedure, both these methods also expose physicians and patients to increased radiation [5, 10].

Real-time intracardiac navigation systems have been developed to address these problems [6, 9]. These systems convert the CT, X-ray and ultrasound images into 3D models of the cardiac anatomy prior to the procedure. Magnetic position tracking coils are embedded into the catheter tip so that graphical models of the catheter can be merged with the 3D heart model and updated in real-time. In this way the use of ionizing radiation during the procedure is significantly reduced. However, significant training and experience is still required to navigate and maintain the catheter tip at a stable position in the work volume, which can be very time consuming [5, 11]. Specifically, difficulty reaching certain positions becomes dominant at high articulation angles, where the device radius of curvature is small and guide lumen forces are high. The flexible and lightweight design of the catheter, in combination with friction inherent in the catheter actuation approach, can result in poor position control fidelity, limiting their application to procedures which do not require precision or resulting in a significant increase in overall procedure time.

To address the dependence on operator skill, reduce training and procedure times, and improve patient outcomes companies are developing robotic systems that use specially designed catheters [6]. To better understand and predict the behavior of these flexible devices, research has been performed on modeling their kinematic relationships using classical approaches [10, 12]. While this has been extremely useful for describing the behavior of many designs, the resulting models are typically static. Design, manufacturing, and variations in the catheter and robot dynamics due to use must be accounted for through periodic calibrations to maintain the design accuracy. Therefore, commercial systems rely on real-time magnetic tracking to provide position feedback. However, position refresh rates of tracking systems range from about 40 to 400 Hz. Averaging is required to increase the accuracy of the system to be clinically useful. However, this can result in a significantly reduced position update intervals. For example, if 200 points are averaged at rate of 200 Hz, a new position measurement is only obtained every 1 s. Other confounding factors, such as sources of interference, from nearby metal objects, other tracking coils, or the CT arm gantry may also degrade tracking accuracy [13]. Automated robotic systems must therefore typically move slowly, when guided by magnetic tracking systems.

The goal of our study was to overcome these limitations by developing a robot that uses adaptive feed-forward control. Empirically learning its inverse kinematics enables the robot to reduce its reliance on feedback from magnetic tracking, while maintaining accuracy and improving its dynamic performance. Initially such a system would use feedback to develop its kinematic models while the physicians guide the catheters manually, or between procedures with the robot practicing on heart models. This would enable the system to work with any available catheter, while automatically compensating for variability in its mechanism. A robust inverse kinematic model would enable the robot to immediately know the actuator positions required to accurately move the catheter tip to a command position. The robot could be instructed to move to a specific location indicated on the 3D heart model, or in the future, even perform a complete procedure in a significantly reduced time by not having to continually correct its position based on the error signal from magnetic tracking. Contributions of this research not only aim to improve the outcomes of patients with AF and others undergoing catheter-based procedures, but also aim to advance the science of intelligent flexible robotics control.

#### II. METHODS

This section can be divided into the design of the robot mechanism (experimental setup), the study of the system dynamics and collection of the test dataset (experimental design) and the design of the machine learning system.

#### A. Experimental Setup

To conduct the learning of the catheter's inverse kinematics, a 3 degree of freedom robotic system was developed, shown in Fig. 2. A standard electrophysiology catheter (J-Type Navi-Star Biosense Webster, Diamond Bar, CA. USA) was mounted in custom designed. 3D printed plastic fixture, linking the catheter with the actuators. Two inlaid thrust roller bearings minimized friction during rotation. Two low-backlash linear actuators (HLD60, Moog Animatics, Milpitas, CA, USA) were used to control the extension  $(m_1)$ and bend  $(m_2)$  of the catheter, with an additional servo motor (2300 Series SmartMotor, Moog Animatics, Milpitas, CA, USA) providing control of the rotation  $(m_3)$ . Each actuator had an RS232 serial interface and incorporated independent proportional-integral-derivation (PID) controllers to achieve a command position within given acceleration and velocity constraints. Three serial to USB bridges enabled independent access to the actuators by the attached robot control computer, running Windows 7 (Microsoft, Redmond, WA) and MATLAB 2014b (Mathworks, Natick, MA).

The catheter was passed through a 1 m long sheath of flexible nylon with a 3 mm ID (5 mm OD) into a acrylic cube (internal length 152 mm). During procedures on patients the sheath is flushed with a saline and anti-coagulant solution, which in addition to lubrication, prevents blood accumulation and clotting in the sheath. The container and sheath were filled with mineral oil to replicate the viscosity of the heart interior and lubrication of the sheath in a non-conductive medium.

The catheter tip was tracked within the cube, by a magnetic tracking system (TrackSTAR<sup>TM</sup> DC EMT, Ascension Technology Corporation). The tracking system was comprised of a control unit that connects to a miniature magnetic sensor (Model 90) which is 7.25 mm in length and 0.9 mm in diameter, and a mid-range magnetic field transmitter. The controller was interfaced to the robot controller via USB. Drivers were provided by the manufacturer to enable MATLAB to make measurements. The sensor was attached with epoxy parallel to, on the circumference, of the elongated catheter body, and aligned with the distal tip. All components were aligned to optimize the accuracy as suggested by Wang et al. [14].



Fig. 2: Experimental setup showing, a) Robotic catheter manipulation mechanism, and b) catheter distal shaft and tracking system in test volume.

#### B. Experimental Design

After the robot was functional we needed to determine whether it was deterministic, by studying the accuracy of the tracking system and repeatability of the catheter kinematics. To determine the tracking system static accuracy, the catheter tip was moved to the "home" position and 2 minutes of position data at 240 Hz was collected. The data was then filtered using a moving average filter with a window size of 0.5 s. The mean of all the smoothed data was used to zero the home position by subtracting it from each window. The resulting data was then uniformly, randomly sampled 500 times, without replacement to obtain the tip x, y and z position.

Since viscoelastic materials are used to construct the catheter, we expected the relationship between the actuator and tip position to exhibit some time-dependency. We decided to minimize the viscoelastic effects by 1) automating the motion profile and measurement intervals to minimize temporal effects, and 2) using the same motion sequence (rotate, extend, bend, relax to home, retract to home and rotate to home) for each experiment. To study the temporal effects and repeatability of catheter tip position, we performed 10 runs, for each run 10 s of position data was collected at "home", then extended the tip by 50 mm ( $m_1$ ) and bent the distal end by moving  $m_2$  by 9 mm, a 5 s pause was performed and the final position was recorded for 10 s and averaged.

To minimize the risk of damaging the robot during development of the learning algorithm it was decided to firstly obtain a comprehensive test dataset. The degree the catheter can bend is a function of its extension. Therefore safe actuator motion ranges needed to be determined to prevent damage. Equation (1) provides the bending versus extension constraint function,  $f_c(\mathbf{m})$ , to prevent excessive stress on the internal tendon causing premature failure. Where,  $\mathbf{m}$  represents a vector of the normalized actuator positions,  $m_1$ ,  $m_2$  and  $m_3$ .

Thereafter, an experimental design was constructed consisting of 3 factors (one for each actuator). Actuator  $m_1$ (extension) had 11 levels (0 to 50 mm in 5 mm steps). Actuator  $m_2$  (bending) had 8 levels (0, 2, 4, 5, 6, 7, 8, 9 mm). Actuator  $m_3$  (rotation) had 12 levels (0 to 330° in 30° increments). After excluding the experiments due to the constraint function, the robot was commanded to perform the remaining 720 experiments in random order. After moving to the command position a 20 s rest period was provided to allow position changes due to viscoelastic effects to subside. Thereafter, 5 s of data was recorded and the mean was subtracted from the home reference position to obtain the tip position. The catheter tip was then returned to its home position and another 10 s interval was provided for relaxation of the materials to occur before the next experiment was performed. This provided a set of real-world data that could be used to test the automatic learning system. Empirically we discovered that each catheter could perform approximately 600 bends before failure of the pull-wire. We managed to collect 538 data points from the catheter under test before it failed, of which 511 were randomly assigned to the training dataset and 27 to the validation dataset to determine the final accuracy of the inverse kinematic model.

$$f_c(\mathbf{m}) = \begin{cases} \infty, & \tanh(m_1 - 0.6m_2^3 - 0.4m_2^2 - 0.3m_2 + 0.45) \ge 0\\ 0, & otherwise \end{cases}$$
(1)

#### C. Machine Learning

In order to learn, the system needed to be able to assess its current knowledge of the real-world and then decide where it needed additional experience. To achieve this, the Model Variance Algorithm (MVA) was used, which uses the deviation in output over a vector of trained models as a measure of uncertainty [15]. A Genetic Algorithm (GA) can then use this uncertainty measure to create hypotheses as to the best actuator positions to test that would reduce the uncertainty in its inverse kinematic model.

A vector of Feedforward Neural Networks (FFNNs) acted as the robots memory of its real-world experience of the forward and inverse kinematic mapping between the three actuator positions  $(m_1, m_2 \text{ and } m_3)$  and catheter tip position  $(x, x_1)$ y and z). A FFNN with one hidden layer was selected since, if given enough neurons in the hidden layer, can fit any finite input-output mapping problem [16]. In addition, a vector of FFNNs can be beneficial as averaging the outputs can improve generalization when trained with noisy or limited data. Hyperbolic tangent sigmoid activation functions were used for the hidden layer, while linear functions were used for the output layer. The Levenberg-Marquardt training algorithm was used, which focuses on minimizing the error at the training data points. Larger standard deviations are therefore found the further away from the point at which a measurement is taken. This provides the MVA metric that can be used to determine whether an arbitrary network output originated from a region near training data or one based on interpolation. Equation (2) provides the MVA score function,  $f_m(\mathbf{m})$ , with, nindicating the number of FFNNs. The sum of  $f_c$  and  $f_m$  is used to evaluate candidate actuator positions.

Ideally, the hypothesized actuator positions would be the next real-world experiment the robot would perform in order to expand its knowledge of its inverse kinematics. For development purposes once the best candidate was selected the real-world experiment with actuator positions associated with the minimum Euclidean norm from the candidate was selected and added to the training dataset. The FFNN was then trained again (without randomizing the weight matrix) and the process repeated. Ideally this loop could continue indefinitely so the robot continues to try and improve its inverse kinematic model to enhance its performance.

$$f_m(\mathbf{m}) = -\overline{\sum_{i=1}^n \sqrt{\frac{\left(net_i(\mathbf{m}) - \overline{net(\mathbf{m})}\right)^2}{n}}}$$
(2)

The MVA technique can also be used to guide the search for the optimal network architecture to maximize generalization (smoothness between training data points). To determine the optimal network size, the vector of neural networks was trained with an increasing number of hidden layer neurons (8 to 50). Training was stopped when the slope of the error reached zero. The mean error at the training data points was use to assess the goodness of fit. The standard deviation of the error using a finer, exhaustive data set with a 1 mm step size was used to assess the generalization. The FFNN exhibiting the minimum inter-sample variance and residual error was selected as the ideal architecture.

#### III. RESULTS

Fig. 3 provides plots of position measurements taken from the catheter tip while stationary in the "home" position. Each marker represents 0.5 s of data representing the mean of 120 position measurements. From these data, ranges for *x*, *y* and *z* were found to be 0.09 mm, 0.09 mm and 0.07 mm (standard deviation  $\pm 0.02$  mm,  $\pm 0.02$  mm, and  $\pm 0.01$  mm), respectively, which is in line with the expected results by Wang et al. [14].

Fig. 4 shows the time-dependent viscoelastic response of the catheter. This results in about 1.0 mm of motion in the z-, and x- axis over a 10 s recording period (start and stop points are 1 s averages). Motion in the x-y plane (not shown) exhibited about a 0.4 mm drift since motion in this axis was only due to the dynamics of the catheter and was not actuated. In spite of these effects good repeatability can be seen.

Fig. 5 provides a scatter plot of the final tip positions due to the experimental design used to build the dataset upon which to study the learning algorithm. The spiral pattern in Fig. 5(a) is due to the position sensor being attached to the catheter tip along its circumference.

Fig. 6 provides the results of the determination of the optimal number of hidden layer neurons. The mean error continues to decrease as the number of neurons increases. This occurs because the FFNN begins to memorize individual data points. However, this typically results in decreased generalization and more uneven interpolation. At the same time the standard deviation can be seen to initially decrease as the FFNN complexity moves closer to the that necessary to fit the training data. However, if the number of neurons increases

further, the standard deviation also increases, suggesting the interpolation smoothness is decreasing. Based on this method 36 hidden layer neurons were selected by taking the minimum of the mean and standard deviation of the error.



Fig. 3: Magnetic tracking system measurements.



Fig. 4: Time dependency of catheter tip position.

Fig. 7(a) shows  $m_1$  and  $m_2$  actuator position with  $m_3 = 0$  and the corresponding mean MVA score,  $f_m(\mathbf{m})$ , for the tip position prediction. Fig. 7(b) shows the cross-section A-A (Fig. 5(a)) after the system has completed 791 learning cycles. The points indicate experiments falling within ±3.0 mm of the section plane. As expected areas closer to data points exhibit a lower MVA score indicating the robot is more certain of the resulting tip position. Areas where the robot could not reach have a high MVA score indicating greater uncertainty.

Once trained, validation was performed with the test dataset. It was found that the model was able to predict the *x*, *y* and *z* tip positions (mean  $\pm 1$  std.) to within  $+0.05 \pm 0.17$  mm,  $-0.21 \pm 0.73$  mm and  $-0.11 \pm 0.62$  mm, respectively.

#### IV. DISCUSSION

Overall, the experiments revealed our methodology resulted in a satisfactory performance of the robot and learning of the inverse kinematic models. Even though only one tip position was obtained for each actuator position, and the pullwire failed before all the experiments could completed, sufficient data was available to learn the inverse kinematics in sufficient detail to allow acceptable positioning accuracy. In future studies a method to prolong the life of the catheter will be needed, or else the reproducibility between catheters needs to be investigated so more data can be collected.



Fig. 5: Final catheter tip positions for experiments showing, a) y-z plane, and b) x-z plane.

During clinical practice the catheter would be firstly located in the sheath at a known reference position in the heart. A 3D model of the heart could then be registered with the patient's anatomy and catheter reference "home" position. Ablation points could be marked on the 3D model and the robot could proceed to perform ablations at the specified locations. For easy clinical investigations, instead of a fully automated robotic system, the robot and inverse kinematic model could also be used for interactive guidance. The system would either allow physicians to perform these procedures remotely, or suggest handle manipulations to move toward a target position. Data from the magnetic tracking system could be used to further optimize the kinematic models in preparation for fully automated procedures. It is conceivable that another method will need to be used to maintain stable contact with the endocardium.

Before clinical applications are explored, there are potential sources of error that must be addressed. A difficultly with feed-forward control is that the effect of all disturbances on the system must be predicted to maintain accuracy. The current kinematic model relates only actuator position to final tip position. Therefore, forces which deform the catheter or distal shaft will result in positioning error. A limitation of this study was that all the experiments were performed under the same conditions and only one measurement was obtained for each run. In practice, the tortuosity of the catheter path will vary, affecting the catheter dynamics and positioning accuracy due to its viscoelasticity and internal friction. There are primarily three sources of disturbances: 1) blood flow, 2) mechanical contact with the endocardium, and 3) friction and viscoelastic effects due to the catheter design. Currently, static 3D models are used therefore attempts are made to minimize disturbances due to blood flow, for example, by creating the 3D heart model at diastole, when blood flow is at a minimum.



Fig. 6: Estimation of optimal FFNN architecture.

The second source of positioning error occurs, when the catheter contacts the endocardium. To improve the robustness of the model to compensate for this error, the model could be extended by incorporating sensors to measure the tip contact force, pull-wire strain and/or actuator forces and torques. In addition, the electrical activity at the tip of the catheter could be measured to detect when it touches the endocardium. Another method would be to use a second robot to control the intracardiac ultrasound imaging catheter, so that real-time visual feedback of the tip position against the heart wall can be collected and used for final positioning guidance.

The third source of positioning error is related to the viscoelasticity of the catheter materials, and the friction between the pull-wire and catheter body, and/or friction between the catheter body and sheath due to deflections along the vascular path to the heart, which may result in variable hysteresis affecting the position accuracy, as was shown in

Fig. 4. This means the relationship between the actuation positions at the handle and the distal tip are nonlinear, timeand path-dependent. Some compensation for these effects could be achieved by incorporating a short history of actuation and elapsed time into the model, and/or by including an estimation of the friction force between then pull-wire and internal catheter lumen by measuring the axial force on the pull-wire. However, this will increase the dimensionality, requiring additional time and measurements to obtain a robust model. It may be better to address these as separate models of correction factors that can be applied to the prediction.



Fig. 7: MVA score with overlay of training data points for: a) actuator position showing constraint function (black line) and uncertainty for tip position and, b) tip position and uncertainty for actuator position through A-A showing training data within  $\pm 3$  mm of section plane.

#### V. CONCLUSION

This research demonstrated successful operation of the automated learning of the robot kinematics. It also showed the feasibility of feed-forward position control for catheter tip positioning. Future work will focus on studying the speed of the robot and reducing the various sources of position error. In addition we are planning to develop both physical and computational 3D artificial heart models that can be used to provide a more realistic environment for the robot to navigate. This will allow us to study techniques and algorithms to enhance the safety of the robot while in the heart, as well as, trajectory planning.

#### VI. ACKNOWLEDGMENTS

We would like to thank the Mayo Foundation for Medical Education and Research for allowing us to use their figure of the cardiac ablation procedure. We would also like to thank Dr Oral and Mr Bloom from the Cardiovascular Center, and Dr Berenfeld from the Center for Arrhythmia Research at the University of Michigan for their assistance.

#### REFERENCES

- Y.-k. Iwasaki, K. Nishida, T. Kato, and S. Nattel, "Atrial Fibrillation Pathophysiology: Implications for Management," *Circulation*, vol. 124, pp. 2264-2274, 2011.
- [2] N. J. Patel, A. Deshmukh, S. Pant, V. Singh, N. Patel, S. Arora, N. Shah, A. Chothani, G. T. Savani, K. Mehta, V. Parikh, A. Rathod, A. O. Badheka, J. Lafferty, M. Kowalski, J. L. Mehta, R. D. Mitrani, J. F. Viles-Gonzalez, and H. Paydak, "Contemporary Trends of Hospitalization for Atrial Fibrillation in the United States, 2000 Through 2010: Implications for Healthcare Planning," *Circulation*, vol. 129, pp. 2371-2379, 2014.
- [3] B. Gorenek and G. Kudaiberdieva, "Cost analysis of radiofrequency catheter ablation for atrial fibrillation," *Int J Cardiol*, vol. 167, pp. 2462-2467.
- [4] A. Ames and W. G. Stevenson, "Catheter Ablation of Atrial Fibrillation," *Circulation*, vol. 113, pp. e666-e668, 2006.
- [5] J. D. Burkhardt and A. Natale, "New Technologies in Atrial Fibrillation Ablation," *Circulation*, vol. 120, pp. 1533-1541, October 13, 2009 2009.
  [6] P. M. Loschak, L. J. Brattain, and R. D. Howe, "Automated pointing of
- [6] P. M. Loschak, L. J. Brattain, and R. D. Howe, "Automated pointing of cardiac imaging catheters," in *Robotics and Automation (ICRA)*, 2013 *IEEE International Conference on*, 2013, pp. 5794-5799.
- [7] J. Jalife, O. Berenfeld, and M. Mansour, *Mother rotors and fibrillatory conduction: a mechanism of atrial fibrillation* vol. 54, 2002.
- [8] A. Noheria, A. Kumar, J. V. Wylie, Jr, and M. E. Josephson, "Catheter ablation vs antiarrhythmic drug therapy for atrial fibrillation: A systematic review," *Arch Intern Med*, vol. 168, pp. 581-586, 2008.
- [9] D. Filgueiras-Rama, A. Estrada, J. Shachar, S. Castrejón, D. Doiny, M. Ortega, E. Gang, and J. L. Merino, "Remote Magnetic Navigation for Accurate, Real-time Catheter Positioning and Ablation in Cardiac Electrophysiology Procedures," *J Vis Exp*, p. 3658, 2013.
- [10] M. Aron and M. M. Desai, "Flexible Robotics," Urologic Clinics of North America, vol. 36, pp. 157-162, 2009.
- [11] Y. Ganji and F. Janabi-Sharifi, "Catheter Kinematics for Intracardiac Navigation," *Biomedical Engineering, IEEE Transactions on*, vol. 56, pp. 621-632, 2009.
- [12] P. Lillaney, C. Caton, A. J. Martin, A. D. Losey, L. Evans, M. Saeed, D. L. Cooke, M. W. Wilson, and S. W. Hetts, "Comparing deflection measurements of a magnetically steerable catheter using optical imaging and MRI," *Med Phys*, vol. 41, p. 022305, 2014.
- [13] Z. Yaniv, E. Wilson, D. Lindisch, and K. Cleary, "Electromagnetic tracking in the clinical environment," *Med Phys*, vol. 36, pp. 876-892, 2009.
- [14] Y. Wang, C. Spangler, B. L. Tai, and A. J. Shih, "Positional accuracy and transmitter orientation of the 3D electromagnetic tracking system," *Measurement Science and Technology*, vol. 24, p. 105105, 2013.
- [15] G. H. Kruger, A. J. Shih, D. G. Hattingh, and T. I. v. Niekerk, "Intelligent machine agent architecture for adaptive control optimization of manufacturing processes," *Adv. Eng. Inform.*, vol. 25, pp. 783-796, 2011.
- [16] K. Hornik, "Approximation capabilities of multilayer feedforward networks," *Neural Networks*, vol. 4, pp. 251-257, 1991.

## Sliding Mode Control for Electromagnetic Levitation System Based on Feedback Linearization

Muhammad Aamir Akram Department of Electronic Engineering The Islamia University of Bahawalpur Bahawalpur, Pakistan <u>amir.akram@iub.edu.pk</u> Ijlal Haider, Habib-Ur-Rehman Khalid MEDICS Lab., Department of Electrical Engineering The University of Lahore Lahore, Pakistan <u>ijlal.haider@ee.uol.edu.pk</u>

Vali Uddin Hamdard University Karachi, Pakistan Vali.Uddin@hamdard.edu.pk

Abstract - Magnetic levitation system has developed into a significant consideration in designing systems which require low losses due to friction and low energy consumption. In this paper the levitation inflexibility for suspending a steel ball and its distinctive application to wipe out friction and build machining accuracy is talked about. The procedures of feedback linearization and sliding mode control are proposed in this paper to accomplish relentless levitation. The procedure of feedback linearization is utilized as a part of this paper to change a nonlinear electromagnetic levitation framework into a linear framework. The proposed controller of sliding mode control is received for enhancing the power and robustness of the electromagnetic levitation system in the vicinity of commotions. Simulation results of sliding mode control demonstrate that the system has more prevailing and powerful bent of stifling disturbance and high inflexibility as equated with that utilizing traditional PID control strategy.

Keywords–Control Systems, Sliding Mode Control, Feedback Linearization, Magnetic Levitation

#### I. INTRODUCTION

Levitation is defined as the stable equilibrium of a body without contact with the earth. The utilization of magnetic power to suspend objects with no obvious method for backing is termed as magnetic levitation or MAGLEV. In view of the rule of operation, magnetic levitation might comprehensively be arranged into two sorts, electrodynamic levitation and electromagnetic levitation [1-4]. This section summarizes the introduction, advantages, applications, fundamental issues, linearization technique and existing SMC controller which is used for MAGLEV system.

With the improvement of control technology, magnetic levitation issues have pulled in numerous scientists and engineers consideration inferable from its friction free dynamic movements. To represent numerous fundamental standards of electrical and electronic engineering a MAGLEV metallic ball framework has been a subject of significant enthusiasm, for example, electromagnetic, circuit design, and control calculations. The MAGLEV framework is nonlinear type and in open loop it is insecure. The system functioning would be vigorously influenced by different disturbances. These testing and intriguing attributes accelerate the design of different controllers keeping in mind the end goal to enhance the system efficiency and performance.

In Magnetic Levitation systems, closed loop consists of actuators, sensors, electromagnet and controllers. Actuator is used as voltage to current convertor which changes over the output voltage from controller into input current for the electromagnet. The explanation behind utilizing this system is to isolated power amplifier from controller. Sensor is used to sense the position of metallic ball which is based on Hall Effect. Electromagnet and a steel ball is our MagLev plant and controller design is our main goal for this paper to enhance the performance and efficiency of MagLev framework.

The objective of the paper is to analyse the MagLev system mathematically and to use the technique of feedback linearization to linearize the system and then design a sliding mode controller to enhance the functioning and efficiency of the framework in the presence of disturbances and uncertainties. So, in the overall procedure of controller design the technique of sliding mode control is used after linearizing the framework by using the technique of feedback linearization.

#### II. LITERATURE SURVEY AND BACKGROUND

The research involved in this paper centrally focused on modelling of system, linearization techniques and designing of controller. In this section the literature review of the work done in the related field is discussed.

Ahmad El Hajjaji and M Ouladsine [1] analysed modelling and nonlinear control of MAGLEV framework. In this paper author proposed a nonlinear model for magnetic levitation system and utilizing this model designed a nonlinear control law in view of differential geometry.

Ximin Shan, Shih-Kang Kuo, Jihua Zhang and Chia Hsiang Manq [5] proposed ultra-precision movement control of a various degrees of freedom magnetic suspension stage. The control architecture designed in this paper is robust nonlinear controller and consists of three parts which are: 1) feedback linearization, 2) force distribution, and 3)  $H_{\infty}$  robust controller for each degree of freedom (DOF) of movement.

Regis Campos Fama, Renato Vilela Lopes and Anderson de Paulo Milhan [6] considered predictive control of a magnetic levitation framework with unequivocal treatment of operational limitations. In this paper the optimal control arrangement is actualized in the receding-horizon strategy, in which the optimization is rehashed at each taking so as to sample moment, by taking into record the new sensor reading.

Subrata Banerjee, Jayanta Pal and Dinkar Parsad [7] considered the execution and efficiency investigation of the controller of an attraction type levitation framework under parametric change. In this work author proposed a cascaded lead compensated controller which consist of an inner current loop and outer position loop and carried out this controller to improve stability and efficiency of single magnet based single axis levitation system.

Shafayet Hossain [8] proposed a pattern of a robust controller for a magnetic levitation system. To overcome nonlinearity problems and achieve stability, author design a robust controller using H-infinity optimization.

Sliding Mode Control (SMC) hypothesis was established and advanced in the previous Soviet Union as a variable structure control framework. SMC is fundamentally an outcome of discontinuous control. SMC was discovered at the beginning of sixties when engineers were looking for ways to design robust control laws. SMC theory initially showed up outside Russia in the mid-1970s when a book by Itkis in 1976 and a review paper by Utkin [13] in 1977 were distributed in English. The SMC reachability condition depends on the hypothesis of Russian mathematician "Lyapunov" on stability of nonlinear systems [9].

### III. SYSTEM DESCRIPTION AND CONTROLLER DESIGN

Magnetic Levitation Systems are finding application in several areas. Fig. 1 shows physical model of MAGLEV system. In order to study MAGLEV System a process is considered which consist of a magnetic ball suspension framework. The object of which is to keep a metal ball suspended in air by adjusting the field strength of an electromagnet. This section describes the basics of system, system dynamics, its modelling and essentials of magnetic levitation system.

Following equations shows the dynamics of electromagnet by using Newton's law of motion:

$$\mathbf{m}\ddot{\boldsymbol{z}} = -\boldsymbol{F}(\boldsymbol{l},\boldsymbol{z}) + \boldsymbol{m}\boldsymbol{g} \tag{1}$$

$$\mathbf{m}\dot{\mathbf{z}} = -k \left[\frac{i(t)}{z(t)}\right]^2 + mg \tag{2}$$

Where

$$F(i, z)$$
 is electromagnet force;

*m* is the mass of the levitated ball;

g denotes acceleration due to gravity;

*z* is the distance of the ball from electromagnet;

*I* is the current through the electromagnet;

- *K* is a coefficient;
- $\mu_0$  is permeability of free space;
- A is the field area of electromagnet core.

Equation (2) describes the complete model of electromagnetic levitation system and indicates that the system is nonlinear due to the  $\left[\frac{i(t)}{z(t)}\right]^2$  term. We have used the technique of feedback linearization to linearize the electro MAGLEV system which is discussed next.



Fig. 1. Physical Model of MagLev System

Feedback linearization is a typical methodology utilized as a part of controlling nonlinear systems. The linear system is acquired through a change of variables and a suitable control input. Feedback linearization [10] may be applied to the nonlinear systems of the form:

$$\dot{x} = f(x) + g(x)u \tag{3}$$

$$y = h(x) \tag{4}$$

The main purpose is to develop a feedback control input

$$\boldsymbol{u} = \boldsymbol{\alpha}(\boldsymbol{x}) + \boldsymbol{\beta}(\boldsymbol{x})\boldsymbol{v} \tag{5}$$

To implement feedback linearization the first step is to find a state-space transformation and to find a state feedback control law. From (2) the states of this system are defined as:

$$x_2 - \dot{z} \tag{7}$$

$$x = [x_1 \, x_2]^T = [z \, z]^T \tag{8}$$

Where  $x_1$  and  $x_2$  are states of the system which are equal to the air gap and rate of change of air gap with respect to time respectively and x in (8) is chosen as the state variable, then the representation of dynamic system in nonlinear state-space equation as:  $x_1 = z = x_2$  (9a)

$$x_2 = \mathbf{Z} = g - \frac{k}{m} \left[ \frac{t}{x_1} \right]^2$$
 (9b)

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -\frac{k}{m} \left(\frac{t}{x_1}\right)^2 + g \end{bmatrix}$$
(10)

According to

$$\boldsymbol{z} = \begin{bmatrix} x_1' \\ x_2' \end{bmatrix} = [f(x) + g(x)\alpha(x) + g(x)\beta(x)\nu]$$
(11a)

and by placing the values of  $\alpha(x)$  and  $\beta(x)$  from Lie Derivative in above equation,

$$\boldsymbol{z} = \begin{bmatrix} x_2 \\ g \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{k}{m} \left(\frac{1}{x_1^2}\right) \end{bmatrix} \begin{bmatrix} \underline{mgx_1^2} \\ k \end{bmatrix} + \left(-\frac{mx_1^2}{k}\right) \end{bmatrix}$$
(11b)

By solving and simplifying the above equation, constants and higher order terms are cancelled out and the obtained linear state space equation is as,

$$\boldsymbol{z} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \boldsymbol{z} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \boldsymbol{v}$$
(12a)

$$y = [1 \ 0]z$$
 (12b)

The system (12a, 12b) is linear and controllable and it can be stabilized by using different types of controllers. For the efficient close loop performance of the system, there is a need of effective and robust controller for this system.

The controller design is subdivided into two steps [9, 11]. These steps include the designing of sliding surface and designing of a control law. Fig. 2 describes the schematic of MagLev system with SMC controller where  $\tilde{z}$  is the desired input signal and z is the actual output of the framework.



Fig.2. Schematic of MagLev System with SMC

The sliding surface for a second order time-invariant framework is a switching line in phase space. It is designed in the form of a line or hyperplane in state-space to compel a sliding motion.

The linearized model is given in (12a) for electromagnetic levitation system. Let the desired output be

$$\mathbf{\vec{z}} = \begin{bmatrix} \mathbf{z}_d & \mathbf{z}_d \end{bmatrix}^T \tag{13}$$

Where  $\mathbf{z}_{d}$  and  $\mathbf{z}_{d}$  are desired levitation height and rate of change of levitation height with respect to time respectively.

The actual output of the electromagnetic levitation system (EMLS) is

$$\mathbf{Z} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{Z}_2 \end{bmatrix}^T \tag{14}$$

The levitation position of the suspended ball is defined as

$$E = \mathbf{Z} - \mathbf{Z} = \begin{bmatrix} \mathbf{g} & \mathbf{d} \end{bmatrix}^T \tag{15}$$

The sliding manifold or surface is formulated as

$$S = CE = C[e \quad e]^T \tag{16}$$

Where S is sliding surface and  $\mathbf{C} - [\mathbf{c}_1 \quad \mathbf{c}_2]$  is the coefficient of sliding surface which is selected so as the rate of change of levitation height with respect to time is automatically adjustable as the disturbances varies and  $\mathbf{c}_1, \mathbf{c}_2$  are selected so as to maintain desired levitation height according to the requirement and to bring the system towards the sliding surface in a finite time.

$$S = \begin{bmatrix} c_1 & c_2 \end{bmatrix} \begin{bmatrix} e & e \end{bmatrix}^T \tag{17}$$

So,

$$S = c_1 e + c_2 e \tag{18}$$

In SMC, the motion is independent of the control. So the control must be designed in such a way that it derives the trajectories towards the sliding surface and maintain these trajectories on the sliding surface once it has been achieved. After the motion of the system enters into the sliding mode, it must satisfy S=0, so the derivative of 'S' is also equal to zero, that is

$$\dot{S} = C\dot{E} = C(\dot{z} - \dot{z}) \tag{19}$$

From (12a)

$$\mathbf{z} = A\mathbf{z} - B\mathbf{v}_{eq} \tag{20}$$

By placing the value of  $\dot{z}$  in (19) and solving it for  $v_{BR}$ 

$$\dot{S} = C\ddot{z} - CAz - CBv_{eq} \tag{21}$$

$$v_{eq} = (CB)^{-1} (C\tilde{z} - CAz)$$
<sup>(22)</sup>

 $\mathcal{V}_{eq}$  is the equivalent control which is used to provide exact sliding surface and brings the system's movement towards the sliding surface.

Reachability condition necessitates that system movement must arrive in finite time at sliding surface. This condition is necessary to be sure that sliding mode starts at some time  $r \ge 0$ , irrespective of the initial state  $\mathcal{X}(0)$ , it is necessary that the state trajectory is always moving towards  $\mathcal{S} = 0$ , whenever  $\mathcal{S}$  is not zero, that is it satisfies  $\mathcal{S} \leq 0$  which called reaching condition. Lyapunov method is used to verify reachability condition [12]. Lyapunov function is stated as follows:

$$\boldsymbol{v} = \frac{1}{2}\boldsymbol{S}^2 \tag{23}$$

Which is positive semi-definite and its derivative to satisfy reachability condition is required as

$$\phi = SS \le 0 \tag{24}$$

A controller must be designed in such a way to satisfy the above reachability condition

 $v \leq 0$  for all values of time. Then the response of the controller must be guaranteed to reach the sliding surface.

By selecting the exponential reaching law as

$$\mathbf{S} = -\mathbf{Ssgn}(\mathbf{S}) - k\mathbf{S} \tag{25}$$

Where the constants are  $\mathcal{E}$ , k and the values of  $\mathcal{E}$ , k > 0. The values of  $\mathcal{E}$ , k are very significant for the behaviour and properties of the system. So the selection of  $\mathcal{E}$ , k should be proper to accelerate the movement and avoid chattering [1]. By solving and substituting equations the control law is obtained as

$$v = (CB)^{-1} (CZ - CAz) + (CB)^{-1} (-\mathcal{E}sgn(S) - kS)$$
(26)

Equation (26) is the addition of two parts,

$$v_{eq} = (CB)^{-1}(CZ - CAZ)$$
(27)

and

$$v_d = (CB)^{-1} (-\mathcal{E}_{\mathcal{F}}gn(S) - kS)$$
(28)

The general form of control law is written as

$$\boldsymbol{v} = \boldsymbol{v}_{eq} + \boldsymbol{v}_{d} \tag{29}$$

Where  $v_{d}$  is the discontinuous part of the control law which ensures the reachability condition, i.e., convergence of system motion to the sliding surface is in a finite time. Combination of equivalent control with switching control should be applied to obtain strong robustness due to the presence of uncertainties and external interferences. The parameters  $\mathcal{E}$ , k,  $c_1$  and  $c_2$  can be properly selected to maintain position error and to achieve desired response. The approximate values of these parameters can be obtained by hit and trial method.

In the vicinity of diverse flaws in switching process, for example, time delays in switching process and little time constants in the actuators, the feedback control law creates a specific dynamic behaviour because of discontinuity in the surface's region, which is normally referred to as chattering. The most regularly referred way to decrease the impacts of chattering has been the alleged piecewise linear or smooth guess of switching component in a boundary layer in the vicinity of sliding surface [3, 4]. The subsequent control with chattering suppression is

$$v = (CB)^{-1} (CZ - CAz) + (CB)^{-1} (-\mathcal{E}sat(S) - kS)$$
(30)

#### IV. SIMULATION AND RESULTS

In this segment, we exhibit some test results demonstrating the object's behaviour in levitation utilizing a proportionalintegral-derivative (PID) controller and sliding mode controller. MATLAB has been adopted in analysis and design of control system for its simplicity and comprehensiveness by researchers. Step by step results are taken by running Simulink blocks.

We have observed the response of both controllers in the presence of disturbance and it is clear that the performance of PID controller is deteriorated as compared to SMC controller. It also shows that linear controller gives good results in safe and protected environments. SMC controller shows much more better and robust performance than conventional PID controller.



Fig. 3. Step Response of SMC Controller



Fig. 4. Step Response of PID Controller



Fig. 5. Response of SMC with Disturbance

#### V. CONCLUSION

In this paper the primary accentuation is given on the technique of feedback linearization to linearize the electro MAGLEV system and sliding mode controller's designing to see the functioning of designed control system in the presence of external disturbances. The functioning of proposed control technique (SMC) is compared with the conventional PID control method. So in the presence of disturbances it is useful to see the functioning of conventional PID controller. The comparison of sliding mode control with PID controller realizes that the system with SMC controller shows not only the rapid response but also disturbance rejection and interference suppression. Parameter tuning techniques can be used to improve the performance of SMC controller and to remove the steady state error of step response as shown in Fig.3.

#### REFERENCES

- Ahmad El Hajjaji and M Ouladsine, "Modelling and nonlinear control of magnetic levitation system," IEEE Transaction on Industrial Electronics, Vol.48, No.4, 2001.
- [2] Walliam G. Hurley and Werner H. Wolfle, "Electromagnetic design of a magnetic suspension system," IEEE Transaction on Education, Vol. 40, No. 2, 1997.
- [3] Faa-jeng Lin, Li-Tao Teng and Po-Huang Shieh, "Intelligent sliding mode control using RBFN for magnetic levitation system," IEEE Transaction on Industrial Electronics, Vol. 54, No. 3, 2007.
- [4] Sung Jun Joo and Jin H. Seo, "Design and analysis of the nonlinear feedback linearizing control for an electromagnetic suspension system," IEEE Transaction on Control System Technology. Vol. 5, No.1, 1997.
- [5] Ximin Shan, Shih-Kang Kuo, Jihua Zhang and Chia-Hsiang Menq, "Ultra Precision Motion Control of a Multiple Degrees of Freedom Magnetic Suspension Stage," IEEE/ASME Transactions on Mechatronics, Vol. 7, No. 1, 2002.
- [6] Régis Campos Fama, Renato Vilela Lopes, Anderson de Paulo Milhan and Roberto Kawakami Harrop Galvão, "Predictive Control of a Magnetic Levitation System With Explicit Treatment of Operational Constraints," 18th International Congress of Mechanical Engineering, November 6-11, 2005.
- [7] Subrata Banerjee, Jayanta Pal, Dinkar Parsad, "Performance study of the controller of an attraction type levitation system under parametric change," Journal of Electrical Systems Vol. 6-3, Pg. 377-394, 2010.

- [8] Shafayet Hossain, Design of a Robust Controller for a Magnetic Levitation System. Wichita State University, Wichita, KS.
- [9] Wilfrid Perruquetti, Jean Pierre Barbot, "Sliding Mode Control in Engineering," Marcel Dekker, Inc New York 2002.
- [10] Liu DE-sheng, Li Jie, Zhang Kun, "The Design of the Nonlinear Suspension Controller for EMS MagLev Train Based on Feedback Linearization," Journal of National University of De Fense Technology, Vol.27, No.2, Pg. 96-101, 2005.
- [11] K. D. Young, V. I. Utkin, and U. Ozguner, "A control engineer's guide to sliding mode control," IEEE Transactions on Control Systems Technology, Vol. 7 No. 3, Pg. 328–342, 1999.
- [12] Zhang T P, "Adaptive fuzzy sliding mode control based on a modified Lyapunov function," Acta Automatica Sinica, Vol.1, No.28, Pg. 137-142, 2002.
- [13] V. I. Utkin, "Variable structure systems with sliding modes," IEEE Transaction on Automatic Control, Vol. 22, No. 2, Pg. 212–222, 1977.
- [14] J. J. Slotine, Applied Nonlinear Control. Englewood Cliffs, NJ: Prentice Hall, 1991.
- [15] Hassan K. Khalil, Nonlinear Systems. NJ 07458: Prentice-Hall, November 2008.

#### Navigation and locomotion of a low-cost Automated Guided Cart\*

Gareth J. Cawood<sup>1</sup> and Igor A. Gorlach<sup>2</sup>

*Abstract*—Automated guided vehicles (AGVs) are widely used for material handling in industry. AGV designs and capabilities vary depending on the complexity of required tasks, the environment and the level of autonomy. As the AGV market is highly competitive, the designer was required to provide a low-cost solution that does not compromise on functionality. An automated guided cart (AGC) is one type of AGV that provides material handling capability with a low level of control and complexity.

This research looks at the development of an AGC for use in assembly lines as requested by a local automotive manufacturer. In particular, this paper focuses on the navigation and locomotion aspects of the AGC, but includes discussion of safety and SCADA related topics. A line-following based method of navigation was implemented making use of a magnetic tape and sensor array. The project is based around a commercial PLC which manages the movement and communication of the AGC. The AGC was tested in a working environment and proved to be capable of material handling tasks as defined by the automotive manufacturer.

#### I. INTRODUCTION

A local automobile manufacturer identified several areas in their factory where they felt automated guided carts (AGCs) could be used effectively. At the time the cost of a commercially available AGC for that purpose was in excess of R 150 000 per unit. Purchasing AGCs at this price was not feasible for the intended implementation. Furthermore, local design firms were also approached but indicated high costs related to development. However, the manufacturer felt that an effective AGC could be designed at a far lower cost. As such a request was made to investigate the design of a low-cost AGC, as well as build a prototype based on those investigations.

Apart from the manufacturer's specific requirements, there was also a general need for low-cost AGCs for commercial use. Although AGCs are finding ever-expanding use in the South African manufacturing industry, there is little local development besides that which takes place at academic institutions. Furthermore current research into AGVs is focused on finding new methods for navigation, mobilisation [1] and odometry [2], not focusing on cost. While important for the future development of AGVs; these systems are costly and difficult to implement into a basic AGC.

The technical requirements for the project were defined as follows: the AGC must

- be able to run 8.5 hours per day, and cover 12 km.
- be able to autonomously navigate a route and stop at markers.
- be capable of towing a load up to 200 kg at a speed of  $1.5 \text{ m} \cdot \text{s}^{-1}$ .
- have dimensions no wider than 400 mm and no higher than 500 mm.
- be able to avoid collisions by detecting potential obstructions.
- be able to be tracked via a SCADA interface.

A low-cost budget of R 35 000 for the AGC was also specified by the manufacturer. This was a major constraint to the project and had to be taken into account throughout the design process.

#### II. DESIGN

The design decisions that would affect the development of the AGC the most related to:

- drive;
- navigation;
- control; and
- sensors.

Each item has several different solutions which impact different aspects of the design. To start the design process several key decisions had to be made regarding the structure and movement of the AGC.

From the start a three-wheeled design was selected. While a four-wheeled solution may provide better stability, it is easy to design around this constraint by manipulating the centre-of-gravity of a three-wheeled platform. Additionally a three-wheeled platform is less likely to be affected by uneven surfaces as all three wheels are continuously on the ground. In contrast, a four-wheeled platform may require suspension to avoid one wheel losing traction when traversing uneven surfaces.

To steer the AGC a differential-drive solution was chosen. The method makes use of two motors driven independently to power two wheels on the same axis. By varying the speed of the motors relative to each other, the direction the AGC faces can be controlled while maintaining forwards motion. This also allows the AGC to rotate about the centre of its axis which increases manoeuvrability. An alternative system was also considered, whereby two wheels are driven by a single motor, and a second motor on the third wheel would control steering. However, the benefits of the differentialdrive system made it the preferred choice.

<sup>\*</sup>This project was made possible thanks to the support of General Motors South Africa through the GMSA Chair of Mechatronics.

<sup>&</sup>lt;sup>1</sup>Gareth Cawood is a master's student at the Mechatronics Department, Nelson Mandela Metropolitan University, Port Elizabeth, South Africa gareth@gcawood.com

<sup>&</sup>lt;sup>2</sup>Igor A. Gorlach is a professor at the Mechatronics Department, Nelson Mandela Metropolitan University, Port Elizabeth, South Africa igor.gorlach@nmmu.ac.za

Another factor that was considered early in the project was how the AGC should be powered. A battery solution was selected over the alternatives (powered rail or induction). Batteries allow for a simpler system with less permanent infrastructure. The specific application required use for only 8 hours with the the remainder of the day available for recharging. Battery systems are used extensively in manufacturing environments in the form of forklifts and tuggers, especially in indoor environments where fumes from gas or diesel powered engines need to be controlled. A battery system has the added benefit of adding mass to what can otherwise be a light vehicle, thus aiding traction.

#### A. Drive

Due to the AGC being powered by batteries, AC motors were immediately discounted for use. Their large size and extra electronics required to allow them to function would negatively impact the space-constrained AGC.

The choice was then between standard DC-, stepper- and servo-motors. Although any motor can be used, standard brushed DC motors were selected. While stepper- and servomotors have good controllability with regards to positional requirements, standard brushed DC motors are more suited to a constant drive situation where smooth movement is beneficial. Coupled with the ease of control and marginal cost benefit, DC motors were selected.

To determine the power requirements for the motor, the maximum total force that the AGC had to overcome was determined. This was done making use of a combination of calculations and physical measurements. The time when the greatest force is required is when the AGC accelerates from a standstill with a trailer attached. The total force was determined to be 192 N. With the selected wheels this equates to a motor torque of 19.2 Nm. Once up to speed this reduces to 10.0 Nm.

A further efficiency factor was included in the motor requirements. A value as low as 65% for certain gear train systems can be evident [3]. This gives a total torque requirement of 29.5 Nm, or 14.8 Nm per motor. Furthermore the motor should be capable of rotating at a rate equal to  $1.5 \text{ m} \cdot \text{s}^{-1}$ . This equates to N = 143.2 RPM.

A motor with torque of 20 Nm and a maximum speed of 230 RPM was selected. At an operating speed of 143 RPM it can deliver 11.4 Nm of torque. Including the 65% efficiency, each motor is required to provide 7.7 Nm of torque at operating speed. These equate to safety factors of 1.36 (startup) and 1.48 (running).

Traction calculations were performed and resulted in a maximum traction force of 115.6 N per wheel. This is sufficient for the 192 N required during acceleration.

While it is possible to control certain motors directly from a PLC or microcontroller, the process is made much simpler, and the level of risk to electrical components is reduced, by using an intermediary motor driver. Furthermore, the selected motors run at 24 V with a running speed current draw of 4 A and a startup current draw of 36 A. These currents would be too high for a PLC or microcontroller to control without a motor controller.

Initial investigations into motor controllers for the automation industry revealed motor controllers costing about R 4 000 per unit. Where two units are required, this is both costly in terms of price and the amount of space the drivers would take up. The best solution that could be found was a *Sabertooth* 2x25 motor driver. While mainly popular amongst hobbyists, there was no indication that the product would not be suitable for this project. A single driver is capable of supplying power to two motors. It is considerably smaller than other drivers and is capable of regenerative braking. Although not locally available, the product could be imported for R 1 600.

#### B. Navigation

Autonomous vehicles have been around since the 1960s and there is constant development and research into new and improved ways to control how they navigate. Historically, forms of line-following have been the most reliable, but several other technologies are available.

Line-following methods require some form of track to be implemented for the AGC to follow. This means that if changes to the route are required, this track has to be physically adjusted. Other technologies permit the reprogramming of the AGC to achieve the same changes.

Dead reckoning assumes the AGC's physical model is known and monitors the commands it makes to motors and actuators to calculate the expected movement. This relies on a number of assumptions, such as knowing the initial starting conditions of the AGC, and that no slipping occurs during movement. With time these calculations are susceptible to drift.

A similar alternative is inertial navigation. An accelerometer on the AGC monitors the forces the AGC undergoes. Again, making use of the initial starting conditions, the AGC can calculate its expected position based on subsequent acceleration readings. Highly accurate accelerometers are required in conjunction with large processing demands.

*Frog AGV Systems* have produced a range of systems based around sensing magnets placed in the floor of the production environment. The magnets are placed in a grid format with 1 m spacings. As the vehicle moves, it constantly locates itself relative to the placement of these magnets. Monitoring the number of magnets it has passed over allows it to keep track of its position on the grid. Both required installation and sensor requirements increase the cost of this system.

Another method for localisation is making use of a system that functions in a similar fashion to the global positioning system (GPS). Whereas GPS makes use of radio signals transmitted by many satellites, the same kind of system can be replicated in a controlled indoor environment, using a combination of radio frequency and ultra-sonic technologies [4].

From the requirements it was decided to use line-following for navigation. While the GPS and magnetic floor methods allow for high accuracy and flexibility, the complexity and expected costs of such systems meant they were not the best suited. Alternatively the dead reckoning and inertial methods, while cheaper, would prove too inaccurate for this application. The added benefit of flexibility was also not important, as the AGC would be used for the same task for an extended time. Line-following has an added reliability factor and low complexity level that makes it advantageous.

Line-following itself can also take several forms. The more advanced systems make use of cameras to differentiate between a line and the surrounding surface. While high levels of accuracy and large resolutions are possible, camera systems require faster, more expensive sensors and processors to interpret the video signals.

*Macome* manufacture a magnetic guidance sensor specifically for AGVs. It is roughly 200 mm wide and gives an analogue output dependent on the sensor's lateral position above a magnetic tape [5]. While used in a similar application by *Xing et al.* [6], and providing good resolution, the cost of R 8 400 per sensor meant that it was too high to meet the budget restraints of the project.

Simpler alternatives lie in the use of either several induction- or infrared (IR) sensors. A minimum of two sensors are required. The sensors are placed next to each other perpendicular to the direction of movement. When the left-hand side sensor is active it means the robot is going off track towards the right and the AGC can be instructed to turn to the left. Likewise when the right-hand side sensor is active, the AGC can be instructed to turn to the right. By using more than two sensors, the resolution can be increased and likewise the level of control.

Induction sensors were chosen to be used. They are capable of sensing the difference between a magnetic and nonmagnetic surface. The floor of a manufacturing environment is generally concrete, so a tape that can be sensed by the induction sensor would have to be used. Induction was chosen over IR (which can differentiate between light and dark surfaces) because it would be more reliable in an industrial setting where a metal track would be less prone to damage. IR sensors can also be affected by sunlight.

An array of five induction sensors were placed at the front of the AGC. While *Makrodimitris et al.* [7] achieved capable navigation with just three sensors, their AGC had a top speed of  $0.32 \text{ m} \cdot \text{s}^{-1}$ , five times slower than the planned speed for the AGC in this project. Five sensors gave a balance between cost and required resolution to navigate effectively. They were spaced at 50 mm intervals to avoid interference.

#### C. Assembly

The manufacturer requested a simple unibody frame. The design needed to be neat and enclosed as much as possible, while still allowing access to critical components. The two lead-acid batteries consumed the most space. They were placed at the centre of the AGC in a cradle for easy replacement if required. The driving wheels were located behind the batteries, giving the AGC slightly increased traction when towing.



Fig. 1. Assembled prototype (sans bumper)

The motors chosen were not able to handle radial forces, so a bearing setup was designed with a shaft running through which connected the motor (with a key) to the wheels. This took the force off of the motors.

A rotating castor wheel was placed at the front of the AGC to support the rest of the weight. Above it a box was constructed to mount the PLC and contain any additional electronic components. On this box additional buttons and switches were mounted for controlling the AGC.

For testing, a perspex plate was formed to hold the induction sensors. It was placed at the front of the AGC. Being at the front allowed the rest of the AGC and the trailer to follow behind the sensors, reducing the likelihood of part of the AGC colliding with something off the track.

The entire frame was welded together, and then powder coated. All components could then be bolted to the frame. Fig. 1 shows the assembled prototype.

#### D. Control

To control the AGC a PLC was selected. While a microcontroller would have been more cost-effective and allowed for slightly more flexibility with programming, a PLC offered several advantages, namely:

- easy integration with automation components (sensors etc.).
- fault finding and reprogramming via computer link.
- known reliability.
- operating voltage aligned with other components.
- no circuit design or interfaces required.

An x86 based computer was also considered, but for many of the same reasons as the microcontroller (as well as additional cost), the PLC solution was chosen.

To test various methods of control for the AGC, a simulation environment was modelled in *Matlab*. When the dimensions of the AGC are known, along with the speed of the motors, one can fully define the state of the AGC.

To simplify the model, a maximum acceleration (as earlier calculated) was programmed in. From this we assume that the robot will always have sufficient torque to react to the models calculated values, and any slipping that occurs is minimal. The environment was modelled using calculations derived by *Lucas* [8].

Initially a 'linear steering' method of control was implemented. This assumed that there were nine different states



Fig. 2. Different states vs active sensors

that the AGC could sense, depending on which sensors were active. While normally only activating one of the five sensors, if the AGC turned sufficiently, two adjacent sensors could be activated at the same time, resulting in four additional states (refer to Fig. 2). A speed was then set for the motors for each of these nine states.

Fig. 3 shows the modelled response of the AGC at low speeds while going round a corner. The first plot is an x-y plane showing the track (solid-line), and the route the AGC took (dotted-line) from its start position (\*). The second plot is the AGC's heading over time, and the third plot indicates which of the five sensors were active at what time.



Fig. 3. Simulated response with linear steering

Thereafter PID control was implemented in the model. While initial tests only using the P multiplier showed results very similar to the 'linear steering' method, implementing either of the other multipliers resulted in an unstable system. This is assumed to be due to the relatively low resolution of the system. Such a finding is supported by the work of *Chang et al.* [9] on motor control as well as a line-following project by *Azlan et al.* [10], which made use of a fuzzy-logic controller. Other types of controllers were not attempted in the model, as the PLC was limited to ladder programming and the built-in PID controller. As such the initial 'linear steering' method was implemented on the AGC in practice.

#### E. SCADA & General

The SCADA system was set up to run independently of the main AGC design. An initial misinterpretation of the PLC's

specifications meant an additional serial port was unavailable for communication. As such an *Arduino* was added for the prototype. An RFID reader and wireless transmitter communicated with the Arduino, which in turn made use of the PLC's IO ports to relay information.

The RFID reader was installed underneath the AGC. This allows it to detect RFID markers that are placed on its track. Each marker has a unique identifier which allows the system to identify at which point in the track the AGC is by what marker it detects. When a valid marker is detected this also triggers an input on the PLC to indicate that the AGC must stop.

Along with the RFID reader, an *XBee* wireless module was included to communicate the position and status of the AGC to a server. A matching *XBee* module is plugged into the server and the data is stored in a *MySQL* database. The system can be expanded to include any number of remote *XBee* modules for use with a system utilising multiple AGCs.

The SCADA overview can be viewed via a web interface on any network connected device with a suitable browser. All information is stored in a *MySQL* database so can be interfaced with existing SCADA systems, or have a unique interface developed for it. Fig. 4 shows one of the possible displays, listing the time, state and location of specific AGCs.

localhost/scada/viewlogs	s.php?nur	ms=18	knumh=10
Home			
Logs			
date	agv	sta	te point
2012-08-04 21:34:19	000	1	04004BEE6E
2012-10-24 22:08:38	001	0	04004BEE69
2012-10-24 22:08:57	001	0	04004BEE69

Fig. 4. SCADA overview showing AGC information

Control of the AGC is done via buttons on the AGC, or via a garage remote-control. From a safety perspective an array of three IR proximity sensors were implemented for collision avoidance. The sensors were capable of sensing objects up to 0.8 m away at a set height of 0.2 m. The sensors are additionally adjustable. While a laser range finder is preferable, at R 10 000 per unit, it did not fit within the budget of the project.

Additionally a bumper was added to the front of the AGC in case an object is not detected by the IR sensors. When the bumper is depressed it triggers an E-stop on the AGC. The AGC will not continue until the bumper is reset.

Fig. 5 presents a schematic overview of how the electronic components were connected.

#### III. TESTING

A specific area at the manufacturer's plant was identified to implement and test the operation of the AGC. It required the AGC to tow a trailer the length of a sub-assembly area (from the end to the beginning). At this point the trailer is manually dehitched, and the AGC returns to the end of the line to collect another trailer. A makeshift tow-hook was bolted to the AGC for the testing.



Fig. 5. Schematic of PLC electronic connections

#### A. Operation

A stainless-steel tape (25 mm wide), with adhesive on one side was used as the track for the AGC. It allowed for easy installation. A route was laid out on the manufacturer's line. Short strips were cut and angled successively to form corners. Small, flat (1 mm) RFID markers were used to indicate stop positions.

During testing the AGC performed one of two kinds of motion, termed three- and five-sensor movement. This relates to how much the AGC would pivot while operating. Three-sensor movement implied that while following a line, the AGC would only pivot slightly  $(8.7^{\circ})$ , never triggering the outer two sensors. Five-sensor movement was characterised by movement that resulted in the outer induction sensors also being activated (17.5°). This is visualised in Fig. 6.



Fig. 6. Difference between 3- and 5-sensor movement

If the AGC started following a straight line it would maintain three-sensor movement. When coming out of a corner the AGC was equally prone to triggering three- or five-sensor movement. Once five-sensor movement started to occur it would not return to three-sensor movement until a further corner was traversed. It was noted that when the trailer is attached, the AGC is 50% less likely to perform fivesensor movement, due to the increased moment of inertia, and resultant decreased rotational acceleration. By reducing the overall speed of the AGC it was also possible to lower the likelihood of the AGC entering five-sensor movement.

#### B. Results

Throughout testing, regardless of whether the AGC performed three- or five-sensor movement, the AGC always maintained its line-following capability. From an aesthetic point of view, five-sensor movement may appear more erratic, but the AGC performs equally well under three- or five-sensor movement, and never lost the track.

During testing the AGC was physically manoeuvred off the track, and in each instance, the AGC would continue for 2.5 s before stopping. It was programmed to react like this when no longer sensing a track for safety reasons.

When accelerating with the trailer hitched, some slipping did occur. This is due to a combination of the surface (painted cement) which is more slippery than what was used in calculations and also the manner in which the motor controller accelerates the wheels: not in a linear fashion over time, but as fast as the system allows. This problem can be addressed by programming with a more appropriate controller and did not impact the performance of the AGC in a measurable way.

The AGC followed a straight line without any problems. However, a second problem occurred when following a curved line while towing the trailer. If a corner was too sharp, the trailer's momentum could push the AGC's back end sideways, causing it to jack-knife. The main reason for this is the large mass of the trailer compared to the AGC. There are several ways to work around this. In the implementation of the AGC the route was designed in a manner to prevent jackknifing. This entailed limiting the radius of corners (greater than 4 m) that the AGC is required to navigate while towing the trailer.

The AGC was able to navigate at the required the speed. At no time during testing did the AGC fail to respond correctly to human input, either via the buttons on the AGC or via the remote control. Furthermore, the IR proximity sensors implemented for collision prevention functioned correctly. Whenever an obstruction was placed in its way, the AGC would come to a stop within 0.1 m of travel after the detection of an obstruction.

According to calculations, the batteries chosen had several times more energy than what the AGC required to run for a full shift. Over three days the AGC was run for 13 hours without requiring a recharge. Although the remaining capacity at this point was not known, 13 hours is already 62% longer than required.

No problems were encountered with the wireless communication or SCADA interface. Whenever a marker was detected, this was successfully communicated to the server along with the state of the AGC at regular time intervals. The parts used in the prototype have a theoretical limit of 40 m line of sight transmission distance. During testing the AGC was never more than 30 m away from the receiving unit. A high power twin of the same product is available with a 1 km range.

One area where the AGC failed to meet design requirements was in identifying markers at which to stop. During three-sensor movement, the AGC only detected an RFID marker with a 94% reliability. During five-sensor movement this dropped to 88% reliability. The main contributor to the problem is the amount of lateral movement of the front of the AGC, this is where the RFID reader was mounted. This resulted in the RFID reader sometimes being too far off the track to detect the marker which was placed on the track.

The final cost of the AGC came to R 26 340 (75% of the allowed budget). The motors, PLC and manufacturing cost of the frame comprised about half the cost of the prototype (R 13 850).

#### C. Recommendations

Research for this project involved the scrutiny of numerous suppliers of sensors and other equipment. Much of the equipment used in this project is designed for use in the industrial environment and as a result had high IP ratings (related to ingress protection) concomitant with high cost. Other components not of as high a standard, for example the IR sensors and motor controller, can be had at a price significantly less than an industrial standard component.

Many of the cheaper versions of required components are not supplied by reputable manufacturers and the quality and reliability of these products is questionable. Finding replacements in the future could become difficult. So while the project came in considerably under-budget, the following recommendations should be considered.

1) Motor Controller: The Sabertooth controller used was not designed with industrial use in mind. Although it did fare well during trials, even during testing in excess of 8 hours, this is minimal compared to the daily, continuous stress the unit will undergo in a production environment. Long term testing is required to determine its reliability. Alternatively more suitable drivers should be selected, although this could compromise the layout and structural design of the AGC.

2) *RFID*: The problem with RFID marker detection can be solved by moving the RFID reader to the back of the AGC. The rear of the AGC deviated far less from the track than the front of the AGC. The problem can be further mitigated by making use of a larger antenna for the RFID reader. Additionally an industrial solution is proposed to replace the barebones electronic solution used in the prototype.

3) *PLC*: Several problems were encountered while trying to implement the PLC into the project. The AGC is fairly

complex and the ladder programming was not sufficient for all desired operations. An alternative controller should be selected which allows for programming via structured text language (STL). Being able to program a more effective control algorithm could also negate the need to change the RFID system. Other topics relating to serial communication can also be addressed in the selection of a different PLC.

#### IV. CONCLUSION

This document highlights the design process for an AGC to be used in the material-handling environment. Support for major design decisions are given and results from testing are presented. The prototype was completed within the specified budget and further met all but one of the requirements. The main problem noted was that the prototype was unable to stop with 100% reliability at designated markers. As such several suggestions were given to solve this problem. In addition recommendations were given to improve the function and general reliability of the AGC as a whole. Beyond the specifications set out by the manufacturer, a SCADA interface with wireless communication was successfully integrated into the AGC allowing for a live overview of a system containing multiple units.

#### REFERENCES

- F. Gomez-Bravo, M.J. Aznar, and G. Carbone, Optimal motion for obstacle stepping in a hybrid wheeled-legged hexapod, in IEEE International Conference on Autonomous Robot Systems and Competitions, Espinho, Portugal, 2014.
- [2] M. Tkocz and K. Janschek, Closed-form metric velocity and landmark distance determination utilizing monocular camera images and IMU data in the presence of gravity, in IEEE International Conference on Autonomous Robot Systems and Competitions, 2014, Espinho, Portugal.
- [3] V. Brusamarello, I. Muller, A. Silva and J. Klein, Low intrusive efficiency estimative of DC motors, in Instrumentation and Measurement Technology Conference, 2013, pp 1255-1260.
- [4] S. Hoehrmann, Entwicklung eines Ultraschall-baserten Ortungssystems fuer Lego Mindstorms Roboter, Research Project, Online (Accessed on 12.3.2012), available at www.informatik.uni-kiel.de/ railway/Downloads/hoehrmann1.pdf, 2004.
- [5] Macome, Macome Guide Sensor, Online (Accessed on 12.3.2012), available at http://www.macome.co.jp/english/gs/gs.htm, 2011.
- [6] W. Xing, L. Peihuang, C. Qixiang, C. Zhou, K. Shen and C. Jin, Design and control of material transport system for automated guided vehicle, in UKACC International Conference on Control, 2012, Cardiff, UK, pp 765-770.
- [7] M. Makrodimitris, A. Nikolajajis and E. Papadopoulos, Semiautonomous color line-following educational robots: Design and implementation, in IEEE/ASME International Conference on Advanced Intelligent Mechatronics, 2011, Budapest, Hungary, pp 1052-1057.
- [8] G. Lucas, A tutorial and elementary model trajectory differential steering the of wheel for system robot 23.6.2013), available actuators, Online (Accessed on at http://rossum.sourceforge.net/papers/DiffSteer/, 2001.
- [9] T. Chang, C. Wang and E. Cohen, Speed control of brushless motor using low resolution sensor, in Proceedings of the American Control Conference, June 2011, pp 25-27.
- [10] N.Z. Azlan, F. Zainudin, H.M. Yusuf, S.F. Toha, S.Z.S. Yusoff and N.H. Osman, Fuzzy logic controlled miniature Lego robot for undergraduate training system, in Second IEEE Conference on Industrial Electronics and Applications, 2007, pp 2184-2188

# Creating a distortion characterisation dataset for visual band cameras using fiducial markers.

Robert Jermy Council for Scientific and Industrial Research Email: rjermy@csir.co.za

*Abstract*—This paper presents two datasets for the purpose of calibrating, evaluating the calibration accuracy, and evaluating stereo vision accuracy of a pair of cameras. The authors provide a baseline that can be used as an initial comparison. This will allow other researchers to perform the same steps and create better algorithms to accurately locate fiducial markers and calibrate cameras. A second dataset that can be used to assess the accuracy of the stereo vision of two calibrated cameras is also provided.

A fiducial marker was presented to a camera in a repeatable sequence of movements at different roll angles. Using the information captured during this stage it is possible to characterise the lens distortion of the camera. The calibrated cameras were then used with the stereo vision assessment dataset to triangulate the displacement between a pair of fiducial markers. The results show that it is indeed possible to use fiducial markers to calibrate visual band cameras but sub-pixel localisation would improve the results further.

All the data collected by the authors have been compiled and is available online at http://prism.csir.co.za/fiducial/.

#### I. INTRODUCTION

This paper investigates the suitability of fiducial markers for the task of distortion characterisation. Fiducial markers are typically used for augmented reality applications [1] due to their characteristics of being both uniquely identifiable and accurately locatable. These same properties make them attractive for camera calibration.

By presenting a known fiducial marker to a camera at a sequence of known positions it is possible to characterise the distortion of the camera. This paper aims to determine if the accuracy of the distortion characterisation is comparable to the method described by de Villiers and Cronje [2].

This paper provides two fiducial marker datasets, one for the calibration and one for the assessment of stereo vision accuracy. An initial analysis of the datasets are also provided. The first dataset allows for the calibration of the stereo cameras' internal parameters and their relative positions. This dataset contains thousands of images of a fiducial marker as it traverses the Field of View (FOV) of two cameras. These images are paired with accurate Three Dimensional (3D) robot arm positions which allows any researchers to perform their own camera calibration. The second dataset, again using fiducial markers, may be used to assess the accuracy of stereo triangulation. The dataset contains synchronised images of a pair of fiducial markers with known displacement which were statically mounted onto a rigid substrate. This will allow Jason de Villiers Council for Scientific and Industrial Research Email: jdvilliers@csir.co.za

other researchers to compare the accuracy of their calibration and stereo vision algorithms by comparing the calculated displacement between the fiducial markers with the known ground truth displacement.

The rest of the paper is organised as follows: Section II provides an overview of the equipment used to capture the datasets. Section III presents and discusses the results. Section IV places the results in context and presents the final findings of the paper.

#### II. EXPERIMENT

This section describes the techniques and equipment used.

#### A. Fiducial Markers

The fiducial markers were generated, identified, and detected using the Chilitags [1] library. The fiducial markers generated by the Chilitags library are all distinct from one another and rotationally asymmetrical. Consequently there is little ambiguity when detecting and identifying a marker that has been presented to a camera.

After the marker has been identified the pixel coordinates of each of the four corners are determined. The usefulness of fiducial markers comes from the fact that they are both uniquely identifiable and accurately locatable. For example in augmented reality applications fiducial markers can be used to insert digital 3D models into live camera images [1].

Four example markers generated using the Chilitags library are shown in Figure 1. Each marker is composed of a square grid of cells and each cell is either black or white. A six by six cell grid containing white and black cells define the fiducial marker pattern. This pattern is surrounded by a black two cell border to provide contrast. The high contrast of the cells means that the markers can be identified even when they subtend a small angle in the FOV of a camera [1].

The datasets provide a full set of images that were captured during the robot movement sequence. When coupled with the accurate robot positions, these datasets will allow interested researchers to characterise the camera pair and compare their triangulation accuracy on the same data used in this paper.

#### B. Equipment

The equipment used for fiducial marker calibration is similar to that described by de Villiers and Cronje [2] and de Villiers et. al [3]. An ABB IRB 120 robotic arm with



(a) Marker with ID 1. (b) Marker with ID 2.



(c) Marker with ID 3. (d) Marker with ID 4.

Fig. 1: Four fiducial markers generated by the Chilitags library.

TABLE I. Cullera specifications.			
Manufacturer	Allied Vision		
Model	GE1600		
Sensitivity Spectrum	400-1000 nm		
Resolution	$1600 \times 1200$		
Pixel size	5.5 µm		
Lens	Schneider	Kowa	
Name	Cinegon 1.8/4.8	LM5JCM	
Nominal focal length	4.8 mm	5.0 mm	
X (mm)	613.97	625.01	
Y (mm)	-97.27	53.64	
Z (mm)	444.66	450.94	
Yaw (deg)	169.53	-178.70	
Pitch (deg)	-0.21	-0.78	
Roll (deg)	0.37	-0.09	

TABLE I: Camera specifications

a stated 10  $\mu$ m accuracy is used to present an identifiable object in a repeatable sequence of known poses to the camera being calibrated. The specifications of the cameras that were evaluated are given in Table I.

The primary difference between this new experiment and previous calibration techniques is that the energy source has been replaced with a fiducial marker. The fiducial marker is mounted on a NewPort M-BK-1A kinematic mount which in turn is mounted at an angle of  $45^{\circ}$  on a NewPort M-BKL-4 kinematic mount, shown in Figure 2, which is mounted on the robotic arm. The roll angles mentioned in this paper are the roll angles of the robotic arm meaning that a roll angle of  $0^{\circ}$  corresponds to a fiducial marker roll angle of approximately  $45^{\circ}$ .

There are some limitations to using a fiducial marker for calibration:

• This technique relies on reflected light for contrast. As such it is most applicable to shorter wavelengths of the optical part of the electromagnetic spectrum i.e.



Fig. 2: Fiducial marker attached to the kinematic mount..

less than 2000 nm.

- Identifying a fiducial marker requires adequate contrast such that the blacks and white cells of the marker are distinct. This means that the calibration must be performed in a well lit environment and the exposure time of the camera increased. This drastically increases the time required to capture the dataset.
- Recognition of the fiducial marker requires sharp focus on the marker. To ensure such focus over the FOV the iris was stopped down, again increasing the dataset capture time.

#### C. Experimental Design

This section serves to explain the full experimental design.

1) Calibration Dataset: A large movement sequence was used to move the robot across the full FOV of the camera. Each movement sequence was composed of a grid made up of 29 rows, 37 columns and two planes. Four full movement sequences were used with each sequence having a different roll of the fiducial marker. The captured positions of two of these completed sequences, overlayed on the camera images, are shown in Figure 3. The coloured blocks in the image represent the average centre of the fiducial marker at each robot position in the movement sequence.

For each position in the robot movement sequence the following steps occur:

- 1) Wait for the robot to reach the specified position.
- Discard current image and next two images from the camera to ensure the exposure period did not include any robot movement.
- 3) Determine over several frames the average position of each corner of the fiducial marker.



(a) Roll angle  $0^{\circ}$  viewed by the Schneider Lens.

(b) Roll angle  $45^{\circ}$  viewed by the Kowa Lens.

Fig. 3: Example distortion grids captured using the fiducial marker with ID 1.

At each robot position the fiducial marker is located, identified, and the pixel positions of the four corners of the marker are located in the image plane. At each robot position multiple images of the fiducial marker are captured and the average corner pixel coordinates are saved to an output file.

Every saved image is  $200 \times 200$  pixels with the fiducial marker centred, except for when the fiducial marker is touching the edges of the image. The output file contains a single line for each captured image which displays the offset from the top left (zero based) of the camera image in pixels as well as the X-, Y-, and Z positions of the robot arm in mm.

2) Assessment Dataset: The assessment dataset makes use of two fiducial markers attached to an aluminium block spaced approximately  $100 \ mm$  apart as shown in Figure 4. The distances between the corners of the markers were measured with a set of Vernier callipers and are shown in Table II. The block was mounted onto the robot and placed in the approximate centre of the FOV of both of the cameras.

The robot was then commanded to move through 10 different roll angles evenly distributed between  $360^{\circ}$ . At each of these roll angles the robot was commanded to five further poses:

- $0^{\circ}$  azimuth and elevation.
- $15^{\circ}$  azimuth and  $0^{\circ}$  elevation.
- $0^{\circ}$  azimuth and  $15^{\circ}$  elevation.



Fig. 4: Pair of fiducial markers used for stereo vision.

TABLE II: Distances between corners of fiducial markers.

Corner Pair	Distance (mm)
Top Left - Top Left	98.97
Top Right - Top Right	98.97
Bottom Right - Bottom Right	99.39
Bottom Left - Bottom Left	99.39

- $-15^{\circ}$  azimuth and  $0^{\circ}$  elevation.
- $0^{\circ}$  azimuth and  $-15^{\circ}$  elevation.

At each robot pose the markers are identified and located. An image of the pair of fiducial markers is captured and the pixel coordinates of the corners of each marker are saved to a file.

Stereo vision techniques were then used to compute the distance between the pair of fiducial markers. The experimental setup for the assessment dataset is shown in Figure 5(b).

#### D. Data Analysis

Only the average positions of corner '1', the top left corner, of the markers was used in this preliminary analysis. Vectors were created using the image centre as the principal point and manufacturers' advertised focal lengths and pixel sizes.

The calibration technique used in this paper is based upon the work of de Villiers and Cronje [2]. The primary change that has been made is that the robot arm mounted energy source has been replaced with a fiducial marker. Brown's model [4], [5] is used to model the distortion of the camera and lens. Three tangential and five radial parameters are fitted to the captured data as described by de Villiers [6] to create an accurate mathematical model of the lens distortion.

The poses of the cameras relative to the robot were determined using vector bundle similarity as per [7].

Triangulation of the markers in the assessment dataset was performed using the distortion calibration and position



(a) Experimental setup for the distortion dataset.

(b) Experimental setup for the assessment dataset.

Fig. 5: Example distortion grids captured using the fiducial marker with ID 1.

characterisation described above to create a ray line in the robot reference system. Then performing the closest point of intersection of these ray lines.

#### III. RESULTS

This section presents the results of the experiment.

#### A. Distortion Calibration Results

Table III shows the results of the distortion calibration as well as position determination at the four different roll angles. The distortion error is the Root of the Mean Square (RMS) deviation of collinear points (in the real world) from the best fit straight line through them. The initial distortion values were 13.767 pixels RMS and 2.779 pixels RMS for the Schneider Cinegon 1.8/4.8 and Kowa LM5JCM respectively. The final position error is the RMS error in degrees between the corresponding vectors in the image based vector bundle and the hypothesised pose vector bundle [7].

These results show that while it is indeed feasible to perform camera calibration using fiducial markers instead of a

		Distortion	Position
Lens	Roll Angle	Error	Error
		(pix. RMS)	(pix. RMS)
	0°	1.06042	0.31117
Schneider	30°	0.99687	0.30186
Cinegon	$45^{\circ}$	0.90994	0.22406
1.8/4.8	70°	0.91149	0.25102
	0°	0.66428	0.24101
Kowa	$30^{\circ}$	0.93907	0.25529
LM5JCM	$45^{\circ}$	0.79221	0.18223
	70°	0.81938	0.19473

TABLE III: Distortion and position results

Light Emitting Diode (LED), it is less accurate than the results reported by de Villiers *et al.* [3]. With errors in the order of one pixel RMS as opposed to less than half a pixel RMS. This degraded performance is due to corner positions reported by the fiducial marker corner detection not being sub-pixel accurate. The characterisations are robust to the angle of roll of the fiducial marker.

#### B. Stereo Vision Results

Table IV presents the stereo vision accuracy results based on the assessment dataset discussed in Section II-C2. It contains the stereo vision accuracy results for the pair of cameras calibrated using the distortion characterisation dataset with the roll angle specified. The robot arm was approximately 200 mm away from the cameras in the X direction.

From Table IV one can see that using the calibration results discussed in Section III-A it is possible to determine the displacement between the pair of fiducial markers to better than 4 mm, which is 2% of the distance of the markers from the camera and 4% the distance between them.

#### IV. CONCLUSION

This paper investigated the feasibility of using fiducial markers for the calibration of visual band cameras. This paper

TABLE IV: Stereo vision measurement of distance between top left corners of the fiducial marker.

Roll	Miss distance (mm RMS)		Displacement error
(deg)	Marker 1	Marker 2	(mm RMS)
0	0.3749	0.3897	1.8628
30	0.3058	0.4214	1.6413
45	0.5528	0.6189	3.7243
70	0.3054	0.4023	3.0830

presents a pair of datasets that can be used to characterise a stereo pair of cameras and then assess their triangulation accuracy.

The same fiducial marker was presented to two cameras by a robotic arm moving through a grid-like movement sequence at varying roll angles. The corners of the fiducial marker were detected using the Chilitags [1] library at each robot position. The position of the marker was captured multiple times at each robot position and the average marker position was stored. This information was then used to characterise the distortion of camera.

A preliminary analysis was performed to verify the correctness of the dataset and the results serve as an initial baseline for the reference of future researchers.

In this assessment it was found that the accuracy of the distortion model created using fiducial markers is slightly worse than that of a camera calibrated using a LED. The roll angles of the fiducial marker did not have a significant effect on the accuracy of the distortion characterisation but did have an effect on the resultant triangulation accuracy.

The preliminary analysis of the assessment dataset was used to verify the accuracy of the camera calibration by verifying the displacement between a pair of fiducial markers mounted onto a rigid aluminium substrate. The accuracy of the stereo vision was found to be within  $3.8 \ mm$  of the ground truth in the worst case, and  $1.7 \ mm$  in the best case.

This dataset will allow other computer vision researchers to use their own methods to locate the fiducial marker, characterise the photogrammetric properties of the camera, determine their triangulation accuracy, and compare their results to those presented in this paper. The dataset is available at http://prism.csir.co.za/fiducial/.

#### A. Future Work

To improve the efficiency of the dataset capture it is possible to mount multiple unique fiducial markers onto a single surface. At every robot position the camera will capture the image location of all the fiducial markers and thus fewer grid points will have to be visited to provide the type of dense grid required for distortion and position characterisation.

The method used to locate the fiducial markers could be improved to provide sub-pixel accurate positions which would further improve the accuracy of the distortion and position characterisation.

#### REFERENCES

- [1] Q. Bonnard, S. Lemaignan, G. Zufferey, A. Mazzei, S. Cuendet, N. Li, A. Özgür, and P. Dillenbourg, "Chilitags 2: Robust fiducial markers for augmented reality and robotics." 2013. [Online]. Available: http://chili.epfl.ch/software
- [2] J. P. de Villiers and J. Cronje, "A method of calibrating a camera and a system therefor," 11 2012, patent Number WO 2014083386 A2.
- [3] J. P. de Villiers, R. S. Jermy, and F. C. Nicolls, "A versatile photogrammetric camera automatic calibration suite for multispectral fusion and optical helmet tracking," in *SPIE Defense, Security and Sensing*, vol. 90860W, 2014, pp. 1–9.
- [4] D. C. Brown, "Decentering distortion of lenses," *Photogrammetric Engineering*, vol. 7, pp. 444–462, 1966.

- [5] —, "Close range camera calibration," *Photogrammetric Engineering*, vol. 8, pp. 855–855, 1971.
- [6] J. P. de Villiers, F. W. Leuschner, and R. Geldenhuys, "Centi-pixel accurate real-time inverse distortion correction," in *Proceedings of the* 2008 International Symposium on Optomechatronic Technologies, ser. ISOT2008, vol. 7266, 2008, pp. 1–8.
- [7] J. P. de Villiers, "Real-time stitching of high resolution video on COTS hardware," in *Proceedings of the 2009 International Symposium on Optomechatronic Technologies*, ser. ISOT2009, vol. 9, 2009, pp. 46–51.

## Topic models for conference session assignment: Organising PRASA 2014(5)

Michael Burke Mobile Intelligent Autonomous Systems Modelling and Digital Science Council for Scientific and Industrial Research South Africa Email: mburke@csir.co.za

Abstract—Conference scheduling and organisation is a particularly laborious task and can be extremely time consuming. While many online conference platforms allow manual topic selection, these can be expensive and typically still require that individual papers be scanned and labelled appropriately before being assigned to reviewers and relevant conference tracks or sessions. This paper shows how the bulk of this process can be automated using topic models. Latent Dirichlet allocation is applied to learn conference topics directly from documents, and a clustering algorithm introduced to separate these into suitably sized conference sessions, determining an appropriate session topic in the process. Conference tracks can then be scheduled by maximising the distance between these session topics, thereby avoiding potential topic conflicts in parallel tracks.

#### I. INTRODUCTION AND RELATED WORK

Conference organisation is particularly time consuming, and the process of allocating papers to sessions can be exhausting, particularly in extremely large, multi-track conferences. Here, conference organisers typically rely on multiple area chairs, each overseeing a particular topic, but this introduces numerous coordination problems, and a reliance on notoriously disorganised academics. Commercial conference scheduling aids like shdlr.com [1] can speed up this process, but generally operate on a drag and drop basis, requiring manual session and track entry.

In an attempt to avoid this laborious task for PRASA, this paper describes an automatic paper allocation approach, which relies on topic models to assign papers to sessions. Our approach operates directly on papers, learns appropriate topics using latent Dirichlet allocation (LDA) [2], and then assigns papers to topics so as to minimise the distance between paper topic distributions within sessions.

Topic modelling is a well established approach to natural language processing that aims to discover themes in large collections of documents. Here, documents are typically modelled as mixtures of topics, each of which contains a vocabulary of words (Figure 1). A number of effective topic modelling techniques have been introduced, most of which extend LDA, the most popular approach to topic modelling. For example, Ramage et al. [3] have proposed a semi-supervised LDA to characterise microblogs and Li and McCallum [4] introduced the Pachinko allocation model, which finds correlations beDeon Sabatta Mobile Intelligent Autonomous Systems Modelling and Digital Science Council for Scientific and Industrial Research South Africa Email: dsabatta@csir.co.za



Fig. 1. Documents are formed by a combination of topics and topics are represented with varying probability in each document.

tween topics in documents using a directed acyclic graph, while Teh et al. [5] propose a hierachical Bayesian model that allows groups of data to be described by coupled Dirichlet processes. Buntine and Mishra [6] have shown that topic modelling using these approaches is both rapid and efficient and can be implemented in parallel. However, despite the numerous extensions to LDA topic modelling, LDA remains ubiquitous across many applications.

Topic modelling techniques are often tested on academic articles or proceedings [7, 8, 9] and provide excellent results suitable for end user applications. For example, JSTOR [10] uses optical character recognition and topic modelling to index documents [11]. However, to the best of our knowledge, topic modelling has yet to be used for conference scheduling, presumably due to the disconnect between the topics found and conference session allocation.

This paper is organised as follows. Section II introduces latent Dirichlet allocation, while Section III describes our approach to conference session assignment along with results when the algorithms are applied to the PRASA 2014 proceedings. Section IV discusses some of the benefits of the proposed approach, and finally, conclusions are presented in Section V.

#### **II. LATENT DIRICHLET ALLOCATION**

Latent Dirichlet allocation [2] is a generative model of documents that is frequently used to find topics from a corpus of documents. Documents are treated as bags of words, drawn from a variety of topics. Let  $\theta_i$  be the distribution of topics in document *i* and  $\phi_k$  the distribution of words for topic *k*. LDA assumes that the *M* documents, each containing  $N_i$  words, in a given collection can be formed using random mixtures over *K* latent topics, with each topic described by a distribution over words.

Topic and word distributions are assumed to arise from Dirichlet distributions parametrised by  $\alpha$  and  $\beta$  respectively,

$$\theta_i \sim \operatorname{Dir}(\alpha),$$
 (1)

$$\phi_k \sim \operatorname{Dir}(\beta).$$
 (2)

The *j*-th word in a corpus is drawn by first choosing a document topic using a categorical distribution with respective event probabilities described by  $\theta_i$ ,

$$Z_{i,j} \sim Cat(\theta_i),\tag{3}$$

and then drawing a word from a second categorical distribution, with event probabilities corresponding to the relevant word distribution for the sampled document topic  $\phi_{z_{i,j}}$ ,

$$W_{i,j} \sim Cat(\phi_{z_{i,j}}). \tag{4}$$

The joint probability of this generative process is

$$p(\mathbf{W}, \mathbf{Z}, \boldsymbol{\theta}, \boldsymbol{\phi}; \alpha, \beta) = \prod_{k=1}^{K} p(\phi_k; \alpha) \prod_{i=1}^{M} p(\theta_i; \beta) \prod_{j=1}^{N_i} p(Z_{i,j} | \theta_i) p(W_{i,j} | \phi_{z_{i,j}},) \quad (5)$$

and the various distributions forming this can be learned using Bayesian inference. Here,  $\mathbf{W}$  is an  $M \times N$  matrix of word identities, with N the total number of words in the corpus vocabulary, and  $\mathbf{Z}$  an N dimensional vector of topics corresponding to each word in the vocabulary.  $\boldsymbol{\theta}$  and  $\boldsymbol{\phi}$  refer to a matrix of document topic distributions and word topic distributions respectively.

Unfortunately, the distributions in (5) cannot be determined in closed form, and a numerical approximation is required. A variational Bayes approximation was used in [12], while Minka and Lafferty [13] apply expectation propagation and Griffiths and Steyvers [14] apply collapsed Gibbs sampling. The latter is used here, as this provides an unbiased estimate of the distributions of interest after an initial burn in period.

#### **III. ASSIGNING DOCUMENTS TO SESSIONS**

The goal of this work is to learn topics directly from papers submitted to a conference and then use these to assign papers to conference sessions. Unfortunately, while LDA is extremely effective at finding topics in a corpus of documents, these topics are not immediately of use in conference session assignment.

 TABLE I

 TOPICS FOUND FOR PRASA/ROBMECH/AFLAT PROCEEDINGS

Topic 1	Topic 2	Topic 3
learn	based	image
training	time	method
programming	model	camera
network	system	pixel
genetic	process	filter
Topic 4	Topic 5	Topic 6
language	feature	robot
speech	hand	control
word	classification	manufacturing
translation	face	system
model	recognition	design

Table I shows the top ranked words that were found for each topic (Somewhat arbitrarily, 6 were used here) when LDA [15] was applied to the PRASA/RobMech/AfLat 2014 proceedings. Training took place on the full corpus of submitted papers (103 papers), but only accepted papers were assigned to sessions. Common English words were removed from the corpus using the collection of stopwords available at [16]. It is clear that the LDA topic model managed to discover the main conference themes; robotics, machine learning, speech and machine translation, computer vision and image processing.

However, papers are mixtures of these topics and as result cannot be assigned directly to a single topic. This is illustrated in Figure 1, which shows the distribution over topics for each paper in the proceedings. While papers could be assigned to the most probable topic for a given document, this could cause errors in the case of application papers, which are typically distributed across multiple topics. In addition, certain topics may contain a larger number of papers. This is common at PRASA, where speech and machine translation papers typically outnumber those in other fields.

Furthermore, there is no guarantee that the number of papers assigned to a given topic in this way will correspond to the required number of conference sessions or tracks. For example, a typical conference consists of three of four 90 minute sessions per day, each comprising 6 papers, and the best papers assigned to each of the 6 topics listed above are highly unlikely to be neatly divisible in this way.

Instead, we allocate papers to sessions so as to minimise the mean of the average sum of distances between the topic distributions of papers assigned to a given session across all sessions,

$$C = \frac{1}{K} \sum_{k=1}^{K} \frac{1}{G_k^2} \sum_{i=1}^{G_k} \sum_{j=1}^{G_k} D(i,j).$$
(6)

Here, K denotes the number of conference sessions,  $G_k$  the number of papers to be presented in session k, and D(i, j) the Euclidean distance between topic probability distributions for papers i and j in session k.

Referring back to Figure 1, it is clear that topic 2 is something of a catch-all topic, describing words that are common across all papers. We prefer to remove this topic when allocating documents to sessions, as it provides very
little information that aids in paper separation, and can result in an undesirable catch-all session. Catch-all topics like this are inevitable in topic modelling, and are easily detected by selecting the topic with the largest sum probability over all documents.

The cost C is minimised in a brute force manner, by initially assigning papers to sessions at random without replacement, and then iterating over all sessions, testing whether swapping a paper in a given session with a paper in another session will result in a reduced cost and exchanging papers if this is the case. This process is repeated a number of times, with different starting points. Algorithm 1 illustrates this more clearly.

Algorithm 1 Paper allocation  $C_{best} = Inf$  $C = C_{best}$ for iter = 1 : MaxIter do for k = 1 : K do  $m_k = G_k$  papers drawn without replacement end for loop  $C_p = C$ for  $a = \operatorname{randperm}(N)$  do for b = 1 to setdiff(1 : N, a) do  $C_{t} = \frac{1}{K} \sum_{k=1}^{K} \frac{1}{G_{k}^{2}} \sum_{i=1}^{G_{k}} \sum_{j=1}^{G_{k}} D(i,j)$ if  $C_{t} < C$  then swap m(a) and m(b) $C = C_t$ else swap m(a) and m(b)end if end for end for if  $C_p = C$  then break loop end if end loop if  $C < C_{best}$  then  $C_{best} = C$  $m_{best} = m$ end if end for

Figure 2 shows the topic distributions when the 2014 PRASA papers are grouped in this manner. Here, we used a schedule of 7 sessions comprising 5 papers, 4 sessions containing 3 papers and a single session of 4 papers, in line with the 2014 PRASA schedule. It is clear that the papers in Figure 2 are successfully clustered by similarity.

After papers have been assigned to sessions, a set of words that best describe each session is desired. Let p(T|D) be the probability of a topic being present given document D, and p(W|T) be the probability of a given word being observed given topic T, both obtained using LDA. The probability of words being observed in a given session can be obtained by first marginalising to find the average distribution over topics for session k,

$$p(T^{k}) = \sum_{i=1}^{G_{k}} p(T_{i}^{k}|D_{i}^{k})p(D_{i}^{k})$$
$$= \frac{1}{G_{k}} \sum_{i=1}^{G_{k}} p(T_{i}^{k}|D_{i}^{k}),$$
(7)

and then using this to determine an appropriate distribution over words for session k,

$$p(W^{k}) = \sum_{j=1}^{K} p(W^{k}|T_{j}^{k})p(T_{j}^{k})$$
$$= \sum_{j=1}^{K} p(W^{k}|T_{j}^{k})\frac{1}{G_{k}}\sum_{i=1}^{G_{k}} p(T_{i}^{k}|D_{i}^{k}).$$
(8)

Selecting a subset of words for which  $p(W^k)$  is greatest provides a set of keywords to describe the session, which can be used to guide the selection of the session title.

While the evaluation of topic model performance is often subjective [17], a comparison between the session assignment found by the proposed approach and that actually used at the conference is useful for performance evaluation. Table II shows paper titles and PRASA 2014 session titles for each session allocation the proposed approach made. It is clear that papers are clustered into suitably similar groupings, but, as expected, the clusters do not quite match those assigned manually. The conference session assignment that was used at the conference is shown in the second column and colour coded according to the manually selected topic a paper was allocated to. Of the automatic paper allocations, almost all seem acceptable, and in many cases groupings appear more sensible that those assigned by hand (Session 4).

The session of most concern is Session 3, which appears to contain somewhat loosely related papers. This grouping appears to have been made because none of these papers seem to fit in a specific topic, and the content appears to be well distributed across all topics. This is clearly exhibited in Figure 2, which shows that the topic probabilities for many of the papers assigned to Session 3 are relatively low. As a result, the algorithm appears to have created an applications session of its own, and assigned 'left-over' papers to this.

Figure 3 shows a trace of the average assignment cost over 30 execution iterations, roughly equivalent to 5 minutes of execution time. Shaded areas indicate standard deviations in cost (the experiment was repeated 100 times). Empirical analysis shows that a cost of less than 0.105 provided reasonable session assignments for this corpus. The results presented above corresponded to an overall assignment cost of 0.094, which was obtained after approximately 8 hours of execution time.

# IV. USEFUL ALGORITHM BY-PRODUCTS

The proposed approach to paper session allocation has some useful by-products. These are briefly discussed below.

TABLE II Papers assigned to sessions

Paper	Session	Title
	Session 1: language speech word translation model	
29	3B: Speech Processing	Number pronunciation in a multilingual environment
32	4B: Natural Language Processing	Unsupervised Topic Modelling on South African Parl
33	5A: Natural Language Processing	An investigation into Spoken Audio Topic Identific
44	5A: Natural Language Processing	Experiments with syllable-based Zulu-English machi
47	4B: Natural Language Processing	Exploring unsupervised word segmentation for machi
	Session 2: robot control manufacturing system design	I C I
5	2B: Robotic Case Studies and Applications	Improvements on a Prosthetic Hand - The UKZN Touch
14	2B: Robotic Case Studies and Applications	The Case for a General Purpose First Response Resc
21	4A: Robotics Sensing and Design	Development of a Two-Wheel Balancing Robot using t
40	4A: Robotics Sensing and Design	Development of an Educational Process Control and
49	4A: Robotics Sensing and Design	Kinematics Analysis and Workspace Investigation of
-	Session 3: image feature method detection camera	, i i i i i i i i i i i i i i i i i i i
1	5B: End User Applications and Systems	Interactive Energy Consumption Visualization
50	5B: End User Applications and Systems	Application of Multi-Objective Local Search to Har
3	4B: Natural Language Processing	SVM Classification of Dr Math Microtext
13	2A: Classifiers AI Machine Learning and Related Topics	Comparison of two detection algorithms for spot tr
24	4A: Robotics Sensing and Design	A vision-based error metric for path following con
	Session 4: learn training programming network genetic	i C
2	5B: End User Applications and Systems	A Novel Approach to Visual Password Schemes Using
8	1A: Natural Language Processing	Context-based Online Policy Instantiation for Mult
20	5A: Natural Language Processing	South African Sign Language Dataset Development an
45	4B: Natural Language Processing	Genetic Programming for Password Cracking. Phase O
48	3A: Image Processing Classifiers and Related Topics	A Comparative Study of Genetic Programming and Gra
	Session 5: image method camera pixel filter	
42	6B: Image Processing	Generation of Super-Resolution Stills from Video
22	3A: Image Processing Classifiers and Related Topics	A study on the sensitivity of photogrammetric came
10	2A: Classifiers AI Machine Learning and Related Topics	Retinal Vessel Segmentation Based on Difference Im
27	2A: Classifiers AI Machine Learning and Related Topics	Automatic infarct planimetry by means of swarm-bas
37	2A: Classifiers AI Machine Learning and Related Topics	A two-Stage Fuzzy c-Means Clustering Algorithm for
	Session 6: language speech word translation model	
25	4B: Natural Language Processing	Comparing Support Vector Machine and Multinomial N
28	5A: Natural Language Processing	Phrase chunking for South African languages: an in
43	1A: Natural Language Processing	Topic Models for Short Text
31	5B: End User Applications and Systems	Performance analysis of a multilingual directory e
36	3B: Speech Processing	Automatic segmentation and clustering of speech us
	Session 7: language speech word translation model	
18	1A: Natural Language Processing	Semi-Supervised Training for Lecture Transcription
35	5A: Natural Language Processing	An English to Xitsonga statistical machine transla
30	3B: Speech Processing	Aligning Audio Samples from the South African Parl
38	3B: Speech Processing	Investigating The Use Of Syllable Acoustic Units F
41	3B: Speech Processing	Effect of Language Resources on Automatic Speech R
	Session 8: robot control manufacturing system design	
4	6A: Robot Design	Towards a Mobile Mapping Robot for Underground Min
9	2B: Robotic Case Studies and Applications	Development of the Electronics Pod for an Underwat
16	2B: Robotic Case Studies and Applications	The Design of a Rugged Low-Cost Man-Packable Urban
	Session 9: feature hand face classification tracking	
7	2A: Classifiers AI Machine Learning and Related Topics	Long-term tracking of multiple interacting pedestr
39	4A: Robotics Sensing and Design	Visual Features in Varying Illumination for Enhanc
46	1B: Image Processing	Single-trial EEG Discrimination between Five Hand
	Session 10: feature image hand face classification	
17	6B: Image Processing	Comparison of background subtraction techniques un
19	1B: Image Processing	Hybrid Age Estimation using Facial Images
34	3A: Image Processing Classifiers and Related Topics	Temporal Classification of FACS AUs using SURF Des
	Session 11: feature hand face classification recognition	
23	3A: Image Processing Classifiers and Related Topics	Automatic classification of sheep behaviour using
26	1B: Image Processing	Vision-based hand motion detection for intent reco
51	6B: Image Processing	Augmenting the L1 Tracker with appearance based tr
	Session 12: robot control manufacturing system design	
6	5B: End User Applications and Systems	CHAMP: a Bespoke Integrated System for Mobile Mani
12	2B: Robotic Case Studies and Applications	Development of a mechatronic transmission control
11	6A: Robot Design	Programmable Logic Control of an Electro- hydrauli
15	oA: Kodot Design	Development of a docking mechanism for self-reconf



Fig. 2. Topic probabilities are similar for documents assigned to respective sessions. Note that the catch-all topic in Figure 1 has been removed here.



Fig. 3. Mean and standard deviation costs for session assignment iterations show typical convergence rates over a 5 minute period.

# A. Scheduling conference tracks

Conference tracks are frequently parallel, and it can be particularly annoying when tracks are poorly scheduled, with talks on similar topics occurring simultaneously. The proposed approach to session allocation can be used to avoid this problem, by scheduling tracks such that distance between session topic distributions,  $P(T^k)$ , is maximised. Figure 4 shows a distance matrix for the sessions selected for PRASA 2014. The figure provides a simple visual aid for track scheduling and allows scheduling conflicts to be avoided. For example, the distance between sessions 1 and 6 (speech and language) is low so these should be scheduled further apart. The session assignment algorithm described in Section III can be used



Fig. 4. A Euclidean distance matrix of the average topic distribution for each session can be used to ensure that session conflicts are avoided.

to do this automatically, in this case by maximising the assignment cost.

# B. Relevance detection

The proposed approach can also be used for relevance detection in the review stages of a conference. Papers that are off topic typically result in a separate latent topic appearing, as they tend to use different vocabularies to relevant topics. As a result, topic distributions for these less relevant papers tend to be peaky, with topic probability mass biased towards a single topic. Thresholding the LDA document topic distributions can flag papers of this type.

# C. Duplicate or similar submission detection

By allocating papers to topics in the proposed manner, papers with similar content tend to be grouped together, making duplicate submission detection easy. This is a particularly useful property for extremely large conferences, where fraudulent submissions can escape the notice of conference organisers.

# V. CONCLUSIONS

This paper has shown how topic distributions learned using latent topic models can be used to assign conference papers to sessions or tracks automatically. Latent Dirichlet allocation was used to find topics in PRASA 2014 proceedings, and the resultant topic distributions used to group papers into sessions. Potential scheduling conflicts are avoided by maxmising the distance between session topic distributions, and the proposed approach allows for relevance and duplicate submission detection. Future work involves scheduling PRASA 2015.

# ACKNOWLEDGMENT

This work was supported by funding from the Council for Scientific and Industrial Research, CSIR, South Africa.

# References

- [1] (2015) shdlr. [Online]. Available: shdlr.com
- [2] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *the Journal of machine Learning research*, vol. 3, pp. 993–1022, 2003.
- [3] D. Ramage, S. T. Dumais, and D. J. Liebling, "Characterizing microblogs with topic models." *ICWSM*, vol. 10, pp. 1–1, 2010.
- [4] W. Li and A. McCallum, "Pachinko allocation: Dagstructured mixture models of topic correlations," in *Proceedings of the 23rd international conference on Machine learning*. ACM, 2006, pp. 577–584.
- [5] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical dirichlet processes," *Journal of the american statistical association*, vol. 101, no. 476, 2006.
- [6] W. Buntine and S. Mishra, "Experiments with nonparametric topic models," in 20th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining, 2014.
- [7] (2014) Nips 2014 papers. [Online]. Available: http: //cs.stanford.edu/people/karpathy/nips2014/
- [8] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, "The author-topic model for authors and documents," in *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, ser. UAI '04. Arlington, Virginia, United States: AUAI Press, 2004, pp. 487–494.
- [9] G. E. Hinton and R. R. Salakhutdinov, "Replicated softmax: an undirected topic model," in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, Eds. Curran Associates, Inc., 2009, pp. 1607–1614.
- [10] (2015) JSTOR. [Online]. Available: www.jstor.org
- [11] D. Blei and J. Lafferty, "Correlated topic models," Advances in neural information processing systems, vol. 18, p. 147, 2006.
- [12] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," in *Advances in neural information processing systems*, 2001, pp. 601–608.
- [13] T. Minka and J. Lafferty, "Expectation-propagation for the generative aspect model," in *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2002, pp. 352–359.
- [14] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences*, vol. 101, no. suppl 1, pp. 5228–5235, 2004.
- [15] A. Riddell, "Ida: 0.3.2," 2014. [Online]. Available: http://dx.doi.org/10.5281/zenodo.12737
- [16] Stopwords. [Online]. Available: https://code.google.com/ p/stop-words/
- [17] J. Chang, S. Gerrish, C. Wang, J. L. Boyd-Graber, and D. M. Blei, "Reading tea leaves: How humans interpret topic models," in *Advances in neural information processing systems*, 2009, pp. 288–296.

# Off-Road Vehicle Active Suspension Using Fluidic Muscle

Phase One - Design and Characterisation

Trevor Stroud

Department of Mechanical Engineering Nelson Mandela Metropolitan University Port Elizabeth, South Africa Trevor.stroud@nmmu.ac.za Schalk Els Department of Mechanical Engineering University of Pretoria Pretoria, South Africa

Hannalie Lombard Department of Mechanical Engineering Nelson Mandela Metropolitan University Port Elizabeth, South Africa

Abstract— Enhancing vehicle stability, while maintaining vehicle control and ride comfort, can significantly improve the chances of accident avoidance in an emergency situation. While traditional vehicle suspension systems provide a compromise between ride and handling, this research study proposes the design of a controllable pneumatic suspension system using fluidic muscle pneumatic actuators, providing a low natural frequency response, for improving ride comfort, but having the ability to respond to severe dynamic maneuvers by increasing the suspension stiffness. In addition, individual corner stiffness and damping will be controlled to manage pitch, roll and yaw movement, as well as ride height, in response to road inputs, direction changes, and vehicle load fluctuations.

The objective of this research is to develop an active suspension control model that will dynamically optimise both the ride and handling parameters of a pneumatic fluidic muscle suspension on a light, off-road experimental motor vehicle. This paper focusses on phase one of the study, namely the research and design of the active pneumatic suspension system, and also provides test results of fluidic muscle characterisation tests to determine their suitability in this application.

*Keywords—active suspension, fluidic muscle, pneumatic suspension, multibody simulation.* 

# I. LITERATURE REVIEW

# A. Ride Comfort and Handling Performance

According to Guglielmino et al. the two main functions of a vehicle suspension system are to isolate the vehicle body from the forces transmitted by external excitation, and to improve road-holding and handling. [1] Generally a softer suspension with low damping will provide better ride comfort due to its ability to absorb road inputs, while a stiffer suspension will provide better handling and road-holding performance, due to its ability to resist pitch and roll movements of the vehicle body, while keeping the tyres in contact with the road surface,

but at the expense of ride comfort. [2] Traditional vehicle suspension designs are therefore usually a compromise between ride comfort and handling requirements, and can be biased towards either end of the ride vs handling spectrum.

In a study by Els [3], vertical acceleration gave the best correlation between subjective and objective ride comfort levels. Fischer and Isermann also considered normalised vertical acceleration as a suitable method of quantifying ride comfort. [4] Vertical acceleration (both magnitude and frequency) of the vehicle body at the driver's position will therefore be the objective variable used in this study to quantify ride comfort.

Gillespie [5] defined the term "handling" as the objective properties of a vehicle when changing direction, including qualities that affect the driver's ability to maintain control of the vehicle. Uys et al. [6] in a search for a single unambiguous objective criterion for evaluating vehicle handling, conducted an extensive study which concluded that roll angle, lateral acceleration and yaw rate were interrelated, and it was suggested that vehicle body roll angle could be a suitable metric for measuring handling performance, if suitable limits were established.

In this study, vehicle handling performance will be evaluated using standardised tests such as the constant radius and ISO severe double lane change test, measuring vehicle body roll, lateral acceleration and yaw rate.

# B. Active Suspension Control Strategies

A popular control strategy focussing on the optimisation of ride comfort, is the skyhook control approach. This approach, initially developed by Crosby and Karnopp [7], is based on the principle that the sprung mass is connected to a hypothetical reference in the sky, which can move in a lateral and longitudinal plane, but limits the vertical movement of the vehicle body. [8] The groundhook control approach is an attempt to improve handling performance of a vehicle by connecting the un-sprung mass, or wheel, to a hypothetical reference on the ground. The objective of groundhook control is to limit the vertical motion of the wheel, to maximise the tyre contact with the road surface. [5]

Robinson [9] developed a control strategy referred to as the "Modified Skyhook with Semi-Active Tracking", which worked with measurements of inertial velocity, relative velocity, and pressure to improve ride comfort.

This study will focus on a hybrid control strategy, which proposes a combination of both skyhook (ride comfort) and groundhook (handling performance) control, where the control law varies between skyhook or groundhook optimisation, depending on selected input variables.

# C. Vehicle Modelling

The purpose of developing mathematical and software simulation models of dynamic systems, is to better understand and eventually predict their behavior. Popular vehicle suspension models range in complexity from simple linear quarter car models to non-linear full vehicle models.

# 1) Quarter Car Vehicle Model

To simplify initial simulations, and gain insight into the basic function of the suspension and control system being modelled, a simplified quarter-car model is often developed which analyses the sprung and un-sprung mass of one corner of the vehicle with two degrees of freedom. This quarter car model does not take into account the pitch and roll motion of the vehicle, and instead focuses on the vertical motion of the body and wheel. [5] The quarter car model consists of all the elements of one corner of the vehicle including a spring and damper element representing the force input into the suspension system. [10] Software simulation programmes such as MSC Adams (Automatic Dynamic Analysis of Mechanical System) have been developed to automatically assemble the equations of motion of the modelled system, and perform numerical integration in order to find the time-domain solutions of these equations. [11]

A quarter car model will be developed in this study to characterise and simulate the components of the suspension system being used, and a test rig will be constructed to test and validate the quarter car model against actual performance.

# 2) Full Vehicle Model

The increasing power of modern computers, has enabled the use of more complex simulation models with higher degrees of freedom, and more realistic results. Els et al. [12] developed a 14 degree of freedom model of a Land Rover Defender vehicle, in a ride vs handling trade-off study with very good results. In this study the quarter car model will be expanded into a full-vehicle model using MSC Adams simulation software.

# 3) Validation of Vehicle Models

According to Bernard and Clover [13], validation is the process of gaining confidence that the calculations yield useful insights into the behaviour of the simulated vehicle. Heydinger et al. [14] stated that in order to validate computer-generated results, they should be compared to vehicle test results. The increasing complexity of computer simulations, provides compounding opportunities for errors in their results, and without a direct correlation with real-life data, these simulations cannot be trusted to represent their physical counterparts. All computer simulation models will be validated using actual test data, where both the quarter car and full vehicle testing will be conducted for this purpose.

# D. Comparison of Coil and Pneumatic Spring Characteristics

# 1) Spring Rate

When comparing a coil spring to a pneumatic spring, the first significant difference occurs when evaluating the force vs displacement characteristics. While the mechanical coil spring has a constant spring rate throughout its stroke (assuming it is wound as a linear spring), a pneumatic spring has a progressive spring rate caused by the physical laws for a polytropic change in state of a gas. The spring rate of an air spring is therefore increasing with an increase in stroke. [15]

# 2) Level Control

While mechanical coil spring suspension systems are designed to be a compromise for most driving situations, gas sprung suspension systems can be easily equipped with levelcontrol, allowing the suspension characteristics to be adjusted to suit fluctuating loads.

# 3) Natural Frequency

In order to provide good isolation from the input side of the suspension, the natural frequency of the output side should be maintained as low as possible (considering the motion sickness limit of 0.5Hz). This means that the spring rate should also be maintained at as low a value as possible. While a mechanical spring rate is constant, the natural frequency will be high at low loads and low at high loads, potentially compromising comfort under low loads. A pneumatic suspension, however is able to provide a constant low level of natural frequency for all load conditions. [15]

# 4) Fluidic Muscle Actuators

Several types of pneumatic actuators, such as cylinders, bellows, pneumatic engines and even pneumatic stepper motors are widely used in industrial and automotive applications, however a less well-known type is that of the so-called Fluidic Muscle. [16]

Fluidic muscles are pneumatic actuators made mainly of a flexible and inflatable membrane, which act like reversebellows, in that they contract on inflation. Their force is not only dependent on the actuation pressure, but also on the state of inflation, providing a second source of spring-like behaviour. [16] They are extremely lightweight, due to the lack of mechanical moving parts, with the core element of the actuator in the form of a membrane. Fluidic muscles can transmit up to ten times the force available from an equivalent-sized cylinder, operating at similar pressures. In addition they employ an hermetically sealed design, providing complete separation between operating medium and the atmosphere, making them ideal for dusty and dirty environments. They are also robust in design, with no opportunity of air leakage through cylinder seals, while having an extremely low internal friction, allowing higher frequency response times when compared to traditional cylinders. [17]

Fluidic muscles have the following dynamic properties:

- A fluidic muscle shortens by increasing its enclosed volume
- A fluidic muscle will contract against a constant load if the pneumatic pressure is increased.
- A fluidic muscle will shorten at a constant pressure if its load is decreased.
- The contraction has an upper limit at which it develops no force and its enclosed volume is maximal.
- For each pair of pressure and load, a fluidic muscle has an equilibrium length.

This behaviour is in absolute contrast to that of a pneumatic cylinder, in that a cylinder develops a force which depends only on the pressure and piston surface area, so that at a constant pressure, it will be constant regardless of the displacement. [16]

Fluidic actuators are contractile devices, generating force and motion in only one direction, and in situations where bidirectional motion is required, two muscles are often coupled together, with each acting in an opposing direction. In this antagonistic setup, one muscle will be moving the load, while the other acts as a brake to stop the load in the desired position. Assuming similar muscles are used, the equilibrium position of the mechanism will be determined by the relative differences in the gauge pressures of the muscles. [16]

Although fluidic muscles have been applied to traditional suspension systems with active control to enhance the vibration performance [18] no applications could be found where the characteristics of fluidic muscles were used alone to actuate an active pneumatic suspension system.

Fluidic muscle actuators are proposed for use as suspension actuators in this study. The selection of these actuators is based on their favourable characteristics, such as light weight, high actuating force, resistance to environmental factors in off-road vehicle applications, fast response time and low internal friction. Two actuations will be mounted antagonistically to control both bump and rebound movement in the suspension, eliminating the need for traditional hydraulic damping.

# II. SUSPENSION DESIGN

The following section summarises key information regarding the lightweight experimental off-road vehicle used in this study, as well as suspension design parameters.

# A. Vehicle Data

The experimental vehicle used in this study was built by Pretoria University according to the 2013 SAE Baja competition rules. Table 1 lists the key specifications of the test vehicle.

#### TABLE I. EXPERIMENTAL VEHICLE SPECIFICATIONS

SAE Baja Experimental Vehicle					
Vehicle Item	Specification				
Overall length	1652 mm				
Overall width	1522 mm				
Overall Height	1500 mm				
Wheelbase	1518 mm				
Vehicle mass with driver	267 kg				
Vehicle mass with driver and test equipment	365 kg				
Weight distribution front/rear	48/52 %				
Front axle mass (including test equipment)	125kg (175 kg)				
Rear axle mass (including test equipment)	142kg (190 kg)				
Briggs & Stratton Engine Power	7.5 kW				
Front suspension travel	130 mm				
Rear suspension travel	120 mm				



Fig. 1. SAE Baja Experimental Vehicle used in this study

# B. Wheel Loads

The wheel loads measured during previous tests on the SAE Baja experimental vehicle are shown in Table 2.

TABLE II. WHEEL LOADS

Wheel average loads at ride height (including data acquisition equipment)						
Front	875 N					
Rear	950 N					
Maximum Wheel Loads (measured on wheel force transducer during dynamic tests)						
Vertical (Fz)	3900 N					
Longitudinal (Fx)	1600 N					
Lateral (Fy)	5200 N					
Moment about Vertical Axis (Mz)	522 Nm					
Moment about horizontal Axis (My)	500 Nm					
Moment about Longitudinal Axis (Mx)	500 Nm					

# C. Active Suspension Control System

The proposed active suspension system architecture will consist of a National Instruments Compact Rio chassis with 40

analogue input and 16 analogue output channels using plug-in modules. Programming is performed in LabVIEW 2014 and Table 3 lists the proposed sensors and actuators to be fitted to the vehicle.

TABLE III.	VEHICLE SENSORS

	Input							
Qty	Sensor	Function						
4	3 Axis Accelerometer 10G	Unsprung mass acceleration (suspension upright)						
3	3 Axis Accelerometer 5g	Sprung mass acceleration (front axle, rear axle, driver position)						
1	3 axis Gyroscope	Sprung mass pitch, roll and yaw motion						
8	Linear rope transducer 300mm	Actuator and wheel displacement						
8	Pressure sensor 0-10 Bar	Pneumatic actuator pressure						
1	Garman GPS Sensor	Vehicle speed & position						
	Ou	tput						
8	Digital Pressure regulator	Fluid muscle actuator pressure						
8	3/2 Way valve	supply pressure isolation						

# III. FLUIDIC MUSCLE CHARACTERISATION

Tensile tests were performed on a Hounsfield tensile testing machine shown in Fig. 2, at various fluidic muscle pre-charge pressures, ranging from 100 kPa to 600 kPa, which is the maximum recommended pressure for the actuator. The test setup included a load cell, 300mm rope transducer and electronic pressure regulator with an analogue pressure output signal. Force, Displacement and Pressure were measured simultaneously on the Hounsfield tensile tester, with data acquisition performed with NI MyRio hardware and LabVIEW software. Both the 40mm and 20mm fluidic muscles were tested to determine their performance characteristics and suitability for use as suspension actuators.

Three types of tests were conducted, namely:

- Force/displacement/pressure at a selected pre-charge pressure at the ride-height displacement
- Force/displacement at a constant pressure using a digital pressure regulator
- Force at a fixed length and varying pressure

From the test data recorded, two types of curves could be drawn up, showing the force vs displacement, and pressure vs displacement, at pre-charge pressures ranging from 100 kPa to 600 kPa. Fig. 3, 4 and 5 show the results of the 40mm muscle tests both at varying and constant pre-charge pressures. The 20mm force vs displacement data is shown in Fig. 6.



Fig. 2. Fluidic Muscle Tests on the Hounsfield Tensile Tester



Fig. 3. 40 x 180 mm Fluidic Muscle Force vs Displacement Curves at Various Pre-charge Pressures



Fig. 4. 40 x 180 mm Fluidic Muscle Pressure vs Displacement Curves at Various Pre-charge Pressures



Fig. 5. 40 x 180 mm Fluidic Muscle Force vs Displacement Curves at Constant Pressure

# IV. SUSPENSION DESIGN OVERVIEW

# A. Front Suspension

# 1) Front Suspension Geometry

Fig. 6 shows the fluidic muscle installation, which is designed to apply tensile forces to the actuator. The active length of the muscle is 180mm, and the installed length between rod eyes is 440mm for the 40mm muscle. A second muscle is mounted diagonally opposite to the bump muscle to perform two functions, namely controlling rebound motion, and to allow the suspension to be preloaded, thereby increasing the effective spring rate.

All Festo fluidic muscles are specified with a 25% maximum contraction, which means that the active muscle length determines the available displacement. The maximum displacement available for both muscles was therefore 25% of 180mm, namely 45mm.



Fig. 6. Front Suspension Kinematics with Fluidic Muscle Installation



Fig. 7. Motion Ratio vs Wheel Travel of Front Suspension

Considering the maximum available wheel travel for the SAE Baja vehicle at 130mm in the front and 120mm at the rear, a motion ration had to be designed to suit the available displacement in the fluidic muscle.

#### 2) Front Syspension kinematic Study

Using the suspension points obtained from the CAD model of the UP Baja Car 17, a three-dimensional model was generated in MSC.ADAMS to analyse the kinematics of the design. The resulting motion ratio shown in Fig. 7 ranged from 2.40 to 2.68 over the full travel of the suspension.

The installation geometry for the rebound (20mm) fluidic muscle produced a motion ratio ranging from 3.11 to 3.00.

# 3) Front Wheel Loads vs Actuator Loads

According to previous testing conducted by the University of Pretoria, the maximum wheel loads did not exceed 3.9 kN. Considering the front actuator motion ratio at full bump travel of 2.68:1, the expected maximum actuator load would be 10.452 kN.

#### B. Rear Suspension

# 1) Rear Suspension Geometry

The same size fluidic muscles were specified for the rear as those used on the front, allowing for more consistent control and characterisation. The resulting motion ratio for the bump muscle ranged from 2.63 to 2.79, and for the rebound muscle, ranged from 4.08 to 3.78. The bump motion ration is therefore similar to the front suspension design.

# 2) Rear Wheel Loads vs Actuator Loads

Considering the rear actuator motion ratio at full bump travel of 2.79:1, and the maximum wheel load of 3.9 kN, the expected maximum actuator load would be 10.881 kN.

# V. CONCLUSION

After completing the fluidic muscle characterisation tests, the following conclusions were drawn regarding the design of the front and rear suspension systems:

*1)* The 20mm muscle did not produce enough force on its own to be useful in this suspension application, unless multiple muscles can be used in parallel, or a single muscle is

6) Evaluate Ride versus Handling Performance

used as an additive force control element in an existing system.

2) The non-linear nature of the 40mm muscle resulted in a relatively low force generated at maximum contraction with a sharp increase closer to the limit of extension.

*3)* The expected wheel loads (both at ride height and under maximum conditions) require actuator forces outside the range of the available force from one 40mm muscle at the required displacement, considering the calculated suspension motion ratios.

4) To accommodate the required actuator forces two 40mm muscles will be required in the bump direction, and one 40mm muscle in the rebound direction.

5) Installing two muscles in parallel allows the option of either controlling them together, or pre-charging one to provide baseline spring characteristics, while the second muscle is controlled to suit the dynamic suspension requirement.

*6)* Packaging of the three 40mm muscles on the test vehicle will be more challenging than two muscles, but can be achieved with careful consideration of the displacement envelope of the respective muscles. A 40 mm muscle expands to 80 mm in diameter when fully inflated.

7) The digital pressure regulator was able to maintain a constant pressure while moving the muscle through the total stoke, providing a useful force vs displacement characteristic.

8) The change in muscle force under constant pressure is produced by the geometry of the muscle reinforcement fibres as the muscle changes in shape through it's displacement.

*9)* This change in force under constant pressure is a characteristic that is unique to fluidic muscles, and is absent on traditional cylinders, which produce a force proportional to pressure.

10) During the fixed length and varying pressure test, the measured force varied greatly from the trend, producing an inconsistent force vs pressure curve. This can be attributed to the mode of operation of the digital pressure regulator which produced abrupt changes in pressure, rather than a gradual increase over time. Further investigation is required to improve the pressure control response in the system.

# VI. NEXT STEPS

The next steps in this study are as follows:

- 1) Manufacture Test Rig
- 2) Model Suspension and Control of Quarter Vehicle
- 3) Rig Test Quarter Vehicle
- 4) Model and Validate Full Vehicle in MSC Adams
- 5) Implement Control System

#### REFERENCES

- Guglielmino, E. Sireteanu, T., Stammers, CW., Ghita,G., Giuclea, M. Semi-active Suspension Control: Improved Vehicle Ride and Road Friendliness. Springer, London. 2008.
- [2] Els, P.S., Theron, N.J., Uys, P.E., Thoresson, M.J. The Ride Comfort vs. Handling Compromise for Off-Road Vehicles. Journal of Terramechanics, Vol.44, pp.303-317. 2007.
- [3] Els, P.S. The Applicability of Ride Comfort Standards to Off-Road Vehicles. Journal of Terramechanics, Vol.42, pp.47-64. 2005.
- [4] Fischer, D., Isermann, R. Mechatronic Semi-Active and Active Vehicle Suspensions, Control Engineering Practice, Vol.12, pp.1353-1367. 2004.
- [5] Gillespie, T.D. Fundamentals of Vehicle Dynamics. Warrendale, PA: Society of Automotive Engineers, Inc. 1992.
- [6] Uys, P.E., Els, P.S., Thoresson, M.J. Criteria for Handling Measurement, Journal of Terramechanics, Vol.43, pp.43-67. 2006
- [7] Crosby, M.J., Karnopp, D.C. The Active Damper. The Shock and Vibration Bulletin, Number 43, Naval Research Laboratory, Washington DC. 1973.
- [8] Harty, D., Branding Vehicle Dynamics. Automotive Engineering International. July 2003, pp.53-60.
- [9] Robinson, W.D. A pneumatic semi-active control methodology for vibration control of air spring based suspension systems. Iowa State University. 2012
- [10] Savaresi, S.M., Poussot-Vassal, C., Spelta, C., Sename, O., Dugard, L. Semi-Active Suspension Control Design for Vehicles. Butterworth-Heinemann; Elsevier, Oxford. 2010.
- [11] Blundell, M., Harty. D. The Multibody Systems Approach to Vehicle Dynamics. Butterworth-Heinemann; Elsevier, Oxford. 2004.
- [12] Els, P.S. The Ride Comfort vs. Handling Compromise for Off-Road Vehicles, PhD Thesis, Department of Mechanical and Aeronautical Engineering, University of Pretoria. 2006. Viewed 17/07/2014 <http://upetd.up.ac.za/thesis/available/etd-07152008-102911/>
- [13] Bernard, J.E., Clover, C.L. Validation of Computer Simulations of Vehicle Dynamics. SAE Tansactions, 940231. 1994.
- [14] Heydinger, G., Garrott, W., Christos, J., Guenther, D. A Methodology for Validating Vehicle Dynamics Simulations. SAE Technical Paper 900128. 1990.
- [15] Bauer, W. Hydropneumatic Suspension Systems. Springer Heidelberg Dordrecht London New York. 2011.
- [16] Daerden, F. Lefeber, D. Pneumatic Artificial Muscles: actuators for robotics and automation. European Journal of Mechanical and Environmental Engineering, 47(1):10–21. 2002.
- [17] Festo product brochure. www.festo.com/rep/en\_corp/assets/pdf/info\_501\_en.pdf. Accessed 05/06/2015.
- [18] Ostasevicius, V. Sapragonas, J. Pilkauskas, K. Staliulionis, D. Suspension based on Artificial Fluidic Muscle for Vibration damping. Kaunas University of Technology. Lithuania. 2012.

# Extracting South African safety and security incident patterns from social media

Vukosi N. Marivate Council for Scientific and Industrial Research vmarivate@csir.co.za

*Abstract*— The use of social media data to gain insights into public phenomena is a potentially powerful tool. We present the use of social media data analysis, connected with crime and public safety incidents, to better understand reoccurring topics and potentially feed into an automated incident detection application. We collected a size-able dataset of Twitter posts (more than 60,000) over a 3 month period by monitoring crime and public safety related keywords linked to accounts. By splitting the data into two categories we are able to extract topics as well as compare and contrast how monitoring official crime and public safety accounts differs from monitoring individuals and organisations that may not be part of that group. Finally we discuss a prototype application, which uses social media data as well as locations to calculate metrics using potential crime and public safety related incidents.

# I. INTRODUCTION

The proliferation of social media services [1], combined with the increase in the number of Internet users all over the world [2], has resulted in a large amount of user created online content. Within this content is a stream of data that can be categorised with different sub-topics that may be of use to researchers or practitioners from different domains. Social media gives one access to a large amount of public thought or reactions that was previously unheard of. As such, a large amount of research has gone into understanding social phenomena, extracting user sentiment, extracting topics as well commercial uses that normally result in products that exploit the social network (collaborative news filtering using social media) [3], [4], [5].

In this paper, we are interested in how South African social media users converse about crime and public safety on social media. Social media in South Africa is increasingly becoming the default communication tool for the public as well as governmental organisations. With this, we have the opportunity to look for patterns within the South African public sphere (strictly people who are active on the internet).

The topic of crime in South Africa is constant due to the high rates of violent crime and perceptions of crime by citizens. Our focus in this area is how we can use social-media data to form a picture of how crime affects individuals in different areas around South Africa. Firstly, this paper delves into understanding the topics that can be extracted from monitoring crime and public safety minded social media accounts. We also look at the interactions other users have with those accounts. In the long term we would like to be able to, as accurately possible, extract and measure crime and public safety incident occurrence to find incident patterns. To be accurate this might take the form of aggregating different pieces of information or sources for the same incident. In such a scheme, more than a single user describing the same incident might give it more credibility. A user also being in the location they mention might increase our confidence. Lastly, as an example might be cross-referencing past incidents we identified with those covered by traditional media and finding patterns that make them more likely to be true.

For the rest of the paper we will go through different aspects of mining social media data. We first present some background on mining social media data for crime and public safety, we then give an overview of the social media space in South Africa. We present the approach we will take in collecting and analysing social media data, present and discuss results. Proceeding that, we present a prototype application that will analyse our data and finally conclude by discussing future work.

# II. MINING SOCIAL MEDIA FOR SAFETY AND SECURITY

There are numerous aspects of mining data from social media [1]. Recent work, in relation to crime and public safety, has focused on collating data from Twitter and crime reports in the City of Chicago to build predictive models for crime using social media [6]. In our work we will be focusing on South Africa, but will not be building predictive models as we do not have access to point data in relation to crime incidents in South Africa. Other work has focused on extracting network data from social media to track terrorism recruitment and communication [7]. We currently are not looking at tracking communication of organised crime networks on social media.

Mining social media for such a use case is obviously not the only way to gather insights into crime and public safety. But, coupled with other sources of information law enforcement as well as governments collect, the insights can have profound influence on mitigating incidents and better responding to them. We believe that, in understanding how incidents occur and how the public reacts and is affected, we can better recommend interventions to mitigate the impact. Whether it be ways to deal with traffic, or finding hotspots of specific types of incidents to more efficiently deploy resources.

#### III. SOUTH AFRICA AND SOCIAL MEDIA

South Africa, as a developing country, is in the unique situation where a large number of internet users are accessing the internet via mobile phones. Couple this with the increasing use of Social Media, we have a potentially large database of data that could be used to gather insights into the South African public. Just like in major countries, government and companies in South Africa use Social Media platforms to communicate as well as a contact point for the public. We use the Social Media platforms to gather insights into South Africans' view into how crime and public safety incidents affect them.

The matter of crime in South Africa became a global debate in the years leading to 2010, when the country hosted the first African World Cup[8], with some studies specifically analysing the effect of crime perceptions on tourist attendance to the games. Especially troubling is the level of violent crime in the country. South Africa has been found to have the same number of violent crimes as the USA and China, countries that are six times and thirty times greater in population, respectively. Amongst the citizens of the country, the dissatisfaction towards safety in the country is not only highlighted through traditional media but it is also discussed in Social Media platforms.

For the South African setting, previous work on analysing social media data in relation to crime had a small Twitter dataset[9]. We go beyond this earlier paper in multiple aspects. At first we present a size-able dataset collected from Twitter over a period of 3 months. Secondly we monitor, not only a few accounts, but also other users mentioning those accounts and extract topics using topic models to better understand what users may be conversing about. Finally we perform topic extraction from the text. We do not tally incidents as there are other considerations to take care of before we can be confident in summarising the number of incidents extracted from Social Media. For us to report tallies at different levels of confidence, we would have to be able to discriminate between posts that have too little information as compared to ones that have an "almost complete" set of information (description of incident, location, time and geolocation from a device).

There are still further opportunities to use data from other, more public, social media platforms such as Facebook. We found that, community focused crime and public safety groups, Community Policing Forums (CPFs), in South Africa, tend to use Facebook as opposed to Twitter. Further we believe there might be more direct communication between members of these forums on messenger platforms such as Whatsapp or WeChat, which are private. For this paper we will only be focusing on public Social Media data from Twitter. The proceeding sections will cover our approach and methods we used to extract insights from the Twitter data.

# IV. APPROACH AND METHODS

We now describe the approach we took to collect data as well as extract patterns from our crime and public safety incident related social media posts.

TABLE I: Twitter account categories and related accounts

Official	Unofficial
@SAPoliceService	@CrimeAirNetwork
@TMPDSafety	@gcalerts
@JMPDSafety	@tps_security
@TrafficSA	@CrimeWatchdog
@CrimeLineZA	@TrafficNewsFeed
@ArriveAlive	@SAcrimefighters
@ER24EMS	@Abramjee

#### A. Data Collection

Data was collected from Twitter for the months of May, June and July 2015. We collected the data by using the Twitter Streaming API with keywords. We monitored Twitter for keywords related to a number of crime and public safety incident related accounts in South Africa. Specifically we monitored accounts that we deemed Official (e.g. Police, Traffic Alerts), or accounts we deemed unofficial (individuals or non-governmental organisations that report or relay messages related to crime and public safety). Examples of a subset of the accounts from both types are shown in Table I.

We monitored posts coming from or being directed at these accounts (mentions). As the data was collected via the Streaming API, we only had access to a small sample of possible tweets [10]. This limitation does not mean we cannot get meaningful results from analysing such a dataset but makes us cognisant that we might miss some of the content.

Focusing on the text in the Twitter posts, we first had to clean the text before being able to use it. To clean the text, we carried out a number of steps. We removed

- English stopwords (including Twitter specific stopwords, e.g. 'rt', 'lrt' etc.),
- special characters,
- URLs,
- words that occur only once,
- and single letter words.

In total we had 35834 posts related to *Official* accounts and 29669 posts related to *Unofficial* accounts. From these posts we can extract patterns, not only about who is posting, but also the topics that may be popular within the network of individuals.

# B. Knowledge Extraction

To extract patterns from the text data, we opted to use a topic extraction approach. We use Latent Dirichlet Allocation (LDA)[11]. In LDA, we treat each social media post as a document d, and each document is comprised of a probabilistic membership to some latent Topics T. Further, each topic, using a bag of words representation, contains a distribution over words W (our vocabulary). As only the words, W, are observable, the goal is to discover the latent topics by using all the documents D and words occurring in those documents. In the end, each document d has a multinomial distribution over the T topics denoted  $\theta_d$ . Each topic t also has a multinomial distribution over the W words  $\phi_t$ . Both  $\theta$ 

and  $\phi$  have Dirichlet priors. There are numerous extensions of the standard LDA model but we will focus on the basic model.

LDA has been applied before to social media data such as Twitter[12], [13]. There are limitations, especially when analysing Twitter data, as the length of each document is short (Twitter posts sizes are limited to 140 characters)[14]. For this work, we will not explore expansion of posts in running LDA on Twitter data. We wish to form a first insight into this type of data and believe the other methods may give us some improvements, but it's best we first understand the limitations as they apply to mining crime and public safety incident data.

# V. RESULTS

In this section, we present the results of analysing our collected Twitter data. We used the Tweepy<sup>1</sup> search API to collect the Twitter data, and used the Natural Language Toolkit (NLTK) [15] and Gensim [16] for Natural Language Processing and LDA respectively. We first will discuss the results of looking at the communication network, afterwards we discuss the extracted topic models.

#### A. Analysis of the Social Network

First, we look at the amount of posts created by different users on each of the datasets. We present the users with the most posts as well as the most mentioned users (other users refer to them). Looking at these characteristics gives us a view of who creates the most content as well as who may be commanding the most attention (being mentioned or talked at the most). There is an opportunity to delve deeper into the network dynamics [17] of the communication network in the dataset. Extracting more patterns, such as network cascades [18] and community detection [19], is beyond the scope of this paper. What we will do is look at the patterns in the outgoing messages (Number of posts) and the incoming messages (Number of mentions)

When looking at the Official category (Figure 1), we note that the highest post frequency, as well as most mentioned users, are traffic related. The accounts @*arrivealive*, @*trafficgp* and @*trafficsa* are primarily concerned with traffic and motor vehicle related incidents. The South African Police Service (@*sapoliceservice*) is the only account, out of the top 4, for both the post frequency as well as mentions, that is not primarily traffic related.

When switching our attention to the Unofficial category (Figure 2, we note that the top posts come from a newspaper (@sonkoerant) and the rest are individuals who normally mention crime related incidents (@abramjee,@pigspotter, @crimeairnetwork, etc.). There are some sprinklings of unrelated accounts (@mariblueruby, @educationgp) that hint at incidents that may have been infrequent. This is highlighted when we compare the post frequency graph to the most mentioned user graph.

In the most mentioned user graph for Unofficial accounts, we see that the top two accounts commanding attention are *abramjee* and @*pigspotter*. @*abramjee* is an account of the head of Crime Line South Africa<sup>2</sup>, who often gets a lot of information directed at him about crimes. @*pigspotter* is an anonymous individual who also commands a lot of attention that is related to crime and traffic. We note that in the Unofficial category, @*sapoliceservice* is mentioned less. This gives us some confidence that the split of the data into these two categories is warranted.

Our analysis of the post and mentions patterns presented an insight into the communication network of the dataset, we now will turn our focus into the topics that can be extracted from dataset's text.

#### B. Analysis of Extracted Topics

To extract topics from the datasets, we ran LDA on our two datasets. We only extracted topics from Twitter posts that had 3 or more words. We wanted to get an insight into how topics extracted from each of the datasets persisted over time. We do this to distinguish topics that, within Twitter users in South Africa, are constant over time from ones that are unique for a short time span. To do this we ran LDA on subsets of the datasets on a week to week basis. Simply, we did not run LDA on full dataset, but on smaller chunks of the data collected for each week. We ran LDA with a topic selection of 10. There could be more topics, but to make the analysis more interpretable, we limited the number of topics. We show a snapshot of the topics extracted for each category, for 3 selected consecutive weeks starting on the 3rd week of May 2015.

First we look at the extracted topics the Official account dataset, Table II. We highlight a sample of safety/security related topics that *persist* over different weeks by using the color blue, while we highlight other safety/security related topics that last only for a week in red. We show only the top 6 words for each topics. What we note is that repeated incidents are of the traffic nature. This makes sense as traffic is a daily occurrence and there may be exacerbating circumstances, such as accidents, that make it worse and also make the public more likely to talk about those circumstances on Twitter. On week 2 we also capture a topic related to Child Protection Week [20], denoted by words {*cpw*,*child*,*minister*,*shabangu*}, which is an event commemorated by the South African Government, and the minister concerned is Minister Susan Shabangu (the minister of Women in South Africa). On that same week we also capture another topic that has to do with shootings at a police station in Alexandra [21] township ({*alexandra*,*killing*,*police*,*people*}).

Moving on to the Unofficial account dataset, Table III, the extracted topics have another set of topics. Topics related to traffic take on different words, with words like *delays* occurring more often. Also extracted are a set of topics to do with events such as a missing person [22] who was unfortunately found murdered (*{ipeleng,moholane,murderedfound}*) or events such as the "End Sealhunt" campaign. Overall, for

<sup>&</sup>lt;sup>1</sup>https://github.com/tweepy/tweepy

<sup>&</sup>lt;sup>2</sup>https://www.crimeline.co.za/



Fig. 1: Network statistics for Official data-set



Fig. 2: Network statistics for Unofficial data-set

	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
	road	thank	stories	durban	n1	pic	traffic	сс	hi	traffic
	rd	good	via	outside	scene	found	lights	n3tour	SS	crash
Week 1	congestion	work	daily	found	lane	dead	road	last	junction	road
	n1	hope	lol	work	north	police	south	crime	team	injured
	n3	n3tour	working	road	accident	allegedly	well	police	repair	accident
	traffic	traffic	cpw	police	people	injured	please	delays	lane	thank
	road	drive	children	driving	police	leaves	road	rd	stationary	cpw
Week 2	near	taxi	road	south	alexandra	crash	police	accident	traffic	drive
	looks	road	minister	roads	thanks	road	man	scene	n1	well
	like	drivers	shabangu	africa	killed	safe	rd	road	ramp	missingchildren
	injured	traffic	news	traffic	injured	thank	traffic	arrested	hi	day
	2	police	taxi	cleared	accident	well	lane	crime	protest	today
Week 3	collision	lights	police	police	leaves	done	n1	good	thanks	buuren
	suspects	still	national	incident	one	gt	crash	please	city	road
	see	working	guys	u	street	hi	rd	work	rob	get

TABLE II: Extracted topics for Official accounts

TABLE IV: Traffic free Official topics for Week 2

Topic	Words
1	alexshooting stopchildabuse cpwshocking details emerge
2	police station safe south killed africa alexandra
3	towards thank arrested durban man near
4	traffic ramp durban still old delays
5	durban woman nice velve attack obviously
6	last seen plates allegedly roodepoort suspects
7	police crime road n3 nope traffic
8	national police drive yet station arrests
9	thanks road closed n3 townhill
10	arrested sapsgp traffic ave injuries suspects

the same period, both data sets have some shared events, such as the Alexandra shooting while others are unique. Traffic seems to be more discussed with Official accounts, linking back to the initial analysis of the most mentioned users in the Official dataset.

To extract more information about topics that may not be related to traffic, in the Official account dataset, we ran LDA once again and remove all documents that primarily fall within the topics of traffic. We do this by calculating the multinomial mixture model of each documents and if the sum of the probability that a document falls into any of the traffic related topics is greater than 0.5, we remove it from the corpus. The resulting topics, without the traffic, are presented in Table IV.

Here we see that the Alexandra shooting incident becomes more apparent (Topics 1 and 2 in Table IV), as well as the topic of what seems like a a Hijacking in Roodepoort (Topic 6 in Table IV). There are still some topics related to traffic (Topics 4, 7 and 9). What this indicates is that there is a need to increase the number of topics for the LDA, which may result in better clusters.

Finally, we also ran LDA with an augmented USER centric model [14]. The USER model expanded the length of each document by aggregating posts for each user and then using that as a document. We augment this model by aggregating USER posts only on a daily basis, not on the complete data. So for a single week, we aggregated the captured tweets per day for each user. The mixture of topics per day may be less likely than mixture of topics from a single user on a weekly or monthly basis. The results, with running this configuration, were still similar to the extracted topics from the normal LDA.

# VI. APPLICATION DEVELOPMENT

Ultimately, the work we are carrying out with the above analysis, is to feed into an automated system that identifies crime and public safety related incidents. The system is envisioned as an information source for different researchers as well as industry practitioners who may be interested in adding to their current tools. Currently in South Africa, crime statistics are only provided in summarised form and as such we may be able to add information about large public incidents. A system diagram of the application is shown in Figure 3. In the system we will take the data from social media, extract topics, rate the information (there is a



Fig. 3: Automated social media crime and public safety detection system

need for a trust model) and use crowd-sourced labelling to build predictive models. In a parallel stream we will extract locations from the social media posts either through tagged locations or location extraction from text.

In future work we will use extracted topic multinomial distributions of new social media posts as features in classifiers. We hope to be able to classify different types of incidents as well as have detectors that can assist us in understanding the content in an automated manner. As shown in the system diagram, we also are working on extracting location information from the social media posts. This will allow us to create metrics that take into account, not only topic or time but also geographic location of incidents. A step from that point would be to develop predictive models for incidents such as in [6] but having some manner to measure trust in the data in such a way as to minimize false positive and false negatives [23].

### VII. CONCLUSION

In this paper, we looked at extracting crime and public safety related topics from data collected from a social media network. We focused on data collected from two different sets of accounts. When analysing the topics, as well as the patterns of who creates most of the content or gets mentioned most, we discover that Traffic related incidents were the most reported on a week to week basis. We are also able to extract topics related to incidents that happened only on a single week we also referred to news articles that referred to those articles. As such, even without complex text expansion, we are able to extract meaningful topics from shortTwitter post data.

There is still a need to extend the number of topics extracted beyond 10, but to do so we will have to find a better way to summarise the extracted topics so that we can communicate the findings. There is still some work that can incorporate supervised LDA, especially so that we can extract known topics and be able to build a more robust topic extraction model for data we have not seen yet. We believe this will move us close in being able to summarise the incidents in such a way that they may be used to augment official South African crime statistics for incidents that happen in public.

Incorporating temporal effects into the topic model will also assist us in better clustered topics, as well as being able to analyse larger amounts of text. We hope to also be able to analyse Twitter posts that are geographically restricted (instead of searching for keywords) and then extracting crime

	Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10
	surprise	africaday	weareafrica	true	oh	like	delays	nkandla	сс	ipeleng
	africaday	happy	sa	back	new	accident	thanks	good	done	moholane
Week 1	yoh	africa	year	road	constructed	awesome	pay	watch	like	19yrs
	au	primary	african	ramp	masjid	child	n1	clean	please	murdered
	anthem	sad	bid	get	indonesia	see	let	zuma	case	found
	good	cape	tragic	cleanup	terrible	hijacked	sad	really	roadblock	end
	police	town	patriotism	saturday	shocking	crashed	emerging	gt	robbed	sealhunt
Week 2	idea	pics	must	delays	damn	vehicle	shooting	upon	road	cruel
	unacceptable	dramatic	duty	source	world	watch	alex	wake	thank	namibia
	great	lets	highway	m2	breaking	recovered	details	farm	done	sealsofnam
	delays	roadblock	rivonia	rwc	thanks	news	good	stabbed	delays	protest
Week 3	source	runaway	congrats	say	cruel	abandoned	denying	african	son	man
	caused	arrested	south	day	end	national	2	memories	inbound	сс
	road	good	via	agree	namibia	found	killers	thank	held	bars
	one	hope	interactors	100	sealhunt	breaking	tronk	one	m4	behind

TABLE III: Extracted topics for Unofficial accounts

and public safety incident topics from that dataset. Currently we have such a dataset, all Tweets from the Gauteng area, that we have been collecting for a month.

#### VIII. ACKNOWLEDGEMENTS

The author would like to acknowledge the contribution of Gary Bezuidenhout, LeeAnne Masilela and Yanga Sisilana to the earlier basic modelling and development of the social media monitoring application.

#### REFERENCES

- J. H. Kietzmann, B. S. Silvestre, I. P. McCarthy, and L. F. Pitt, "Unpacking the social media phenomenon: towards a research agenda," *Journal of Public Affairs*, vol. 12, no. 2, pp. 109–119, 2012.
- [2] W. Bold and W. Davidson, "Mobile broadband: redefining internet access and empowering individuals," *The Global Information Technology Report 2012: Living in a Hyperconnected World*, 2012.
- [3] S. Staab, P. Domingos, J. Golbeck, L. Ding, T. Finin, A. Joshi, and A. Nowak, "Social networks applied," *Intelligent Systems, IEEE*, vol. 20, no. 1, pp. 80–93, 2005.
- [4] A. M. Kaplan and M. Haenlein, "Users of the world, unite! the challenges and opportunities of social media," *Business horizons*, vol. 53, no. 1, pp. 59–68, 2010.
- [5] P. J. Carrington, J. Scott, and S. Wasserman, *Models and methods in social network analysis*, vol. 28. Cambridge university press, 2005.
- [6] M. S. Gerber, "Predicting crime using twitter and kernel density estimation," *Decision Support Systems*, vol. 61, pp. 115–125, 2014.
- [7] G. Dean, P. Bell, and J. Newman, "The dark side of social media: review of online terrorism," *Pakistan Journal of Criminology*, vol. 3, no. 3, pp. 103–122, 2012.
- [8] R. George and K. Swart, "International tourists' perceptions of crimerisk and their future travel intentions during the 2010 fifa world cup<sup>TM</sup> in south africa," *Journal of Sport and Tourism*, vol. 17, no. 3, pp. 201– 223, 2012.
- [9] C. Featherstone, "The relevance of social media as it applies in south africa to crime prediction," in *IST-Africa Conference and Exhibition* (*IST-Africa*), 2013, pp. 1–7, IEEE, 2013.
- [10] F. Morstatter, J. Pfeffer, H. Liu, and K. M. Carley, "Is the sample good enough? comparing data from twitter's streaming api with twitter's firehose," arXiv preprint arXiv:1306.5204, 2013.
- [11] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," the Journal of machine Learning research, vol. 3, pp. 993–1022, 2003.
- [12] W. X. Zhao, J. Jiang, J. Weng, J. He, E.-P. Lim, H. Yan, and X. Li, "Comparing twitter and traditional media using topic models," in *Advances in Information Retrieval*, pp. 338–349, Springer, 2011.
- [13] J. Mazarura, A. de Waal, F. Kanfer, and S. Millard, "Topic modelling for short text," *Proceedings of the Pattern Recognision Association of South Africa (PRASA 2014)*, 2014.
- [14] L. Hong and B. D. Davison, "Empirical study of topic modeling in twitter," in *Proceedings of the First Workshop on Social Media Analytics*, SOMA '10, (New York, NY, USA), pp. 80–88, ACM, 2010.

- [15] S. Bird, "Nltk: the natural language toolkit," in *Proceedings of the COLING/ACL on Interactive presentation sessions*, pp. 69–72, Association for Computational Linguistics, 2006.
- [16] R. vRehuuvrek and P. Sojka, "Gensim-python framework for vector space mo delling," *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, 2011.
- [17] D. Easley and J. Kleinberg, Networks, crowds, and markets: Reasoning about a highly connected world. Cambridge University Press, 2010.
- [18] J. Leskovec, M. McGlohon, C. Faloutsos, N. S. Glance, and M. Hurst, "Patterns of cascading behavior in large blog graphs.," in *SDM*, vol. 7, pp. 551–556, SIAM, 2007.
- [19] D. Chakrabarti and C. Faloutsos, "Graph mining: Laws, generators, and algorithms," ACM Computing Surveys (CSUR), vol. 38, no. 1, p. 2, 2006.
- [20] M. Mortlock, "Calls for increased child safety as child protection week draws to close," *Eye Witness News*, June 2015.
- [21] EWN, "5 dead after alex police station shooting," Eye Witness News, June 2015.
- [22] T. Ntobela, "Missing ipeleng moholane found dead in Midrand," *The Citizen Newspaper*, May 2015.
- [23] K. Starbird, J. Maddock, M. Orand, P. Achterman, and R. M. Mason, "Rumors, false flags, and digital vigilantes: Misinformation on twitter after the 2013 boston marathon bombing," *iConference 2014 Proceedings*, 2014.

# Using generalized maxout networks and phoneme mapping for low resource ASR- a case study on Flemish-Afrikaans

Reza Sahraeian<sup>1</sup>, Dirk Van Compernolle<sup>1</sup> and Febe de Wet<sup>2</sup>

Abstract—Recently, multilingual deep neural networks (DNNs) have been successfully used to improve under-resourced speech recognizers. Common approaches use either a merged universal phoneme set based on the International Phonetic Alphabet (IPA) or a language specific phoneme set to train a multilingual DNN. In this paper, we investigate the effect of both knowledge-based and data-driven phoneme mapping on the multilingual DNN and its application to an under-resourced language. For the data-driven phoneme mapping we propose to use an approximation of Kullback Leibler Divergence (KLD) to generate a confusion matrix and find the best matching phonemes of the target language for each individual phoneme in the donor language. Moreover, we explore the use of recently proposed generalized maxout network in both multilingual and low resource monolingual scenarios. We evaluate the proposed phoneme mappings on a phoneme recognition task with both HMM/GMM and DNN systems with generalized maxout architecture where Flemish and Afrikaans are used as donor and under-resourced target languages respectively.

*Index Terms*—Low resource ASR, phoneme mapping, Kullback Leibler Divergence, multilingual deep neural network.

# I. INTRODUCTION

Exploiting out-of-language data to develop high performance speech processing systems for low-resource languages has been extensively used recently [1][2]. However, sharing the knowledge across various languages is not a straightforward task because of differences such as different sets of subword units. In the literature, a common approach towards this is the creation of a universal phoneme set by first pooling the phoneme sets of different languages together and then merging them based on their similarity in both knowledge-based and data-driven fashions [3][4]. Knowledge-based phoneme mapping needs prior expertknowledge of a phonetician and is an appropriate approach when we have no data for the target language. In practice, however, we usually have at least a few hours of data. To benefit from the available data, data-driven phoneme mapping can be used instead [5][6].

In the realm of multilingual neural networks [7], creating the target phoneme set for the multilingual training is commonly done (a) by joining of language-specific phoneme sets, (b) training neural networks where each language has its own output layer or (c) by mapping to a global phoneme set. The first two approaches have been successfully used when sufficient amount of training data for each language is available [8][9]. In the case of limited training data, however, using information from high resource language(s) by merging phoneme sets may be beneficial [10]. While the common approach for multilingual DNN training is that each language has its own output layer, our goal is to investigate if better performance can be gained by knowledge-based and data-driven phoneme mapping and which one performs best. This is a tricky issue as it depends on the languages. For example, if two languages are closely-related, IPA based mapping may work sufficiently well. Thus, in this paper, we conduct a case study for two related languages: Flemish and Afrikaans [12].

The data-driven approach we used is based on learning a phoneme mapping table by calculating KLD between pairs of phonemes in Flemish and Afrikaans. It is worth noting that similar works exist where a data-driven phoneme mapping is addressed by making the confusion matrix using multilingual neural networks [13][11]. However, the reported performance mostly degrades compared to the knowledge-based method. Moreover, there are two aspects in which this paper differs from [13]. First, the latter dealt with languages with moderate amounts of data and therefore DNN training where each language has its own output layer yields the best results; whereas, we deal with the resource-scarce target language and phoneme mapping is beneficial. Moreover, our approach is more flexible as we may assign more than one phoneme from Afrikaans to each phoneme of Flemish based on the confusion scores.

In addition, deep maxout networks have achieved improvements in various aspects of acoustic modelling for large vocabulary speech recognition systems including underresourced and multilingual scenarios [14][15]. In this paper, we investigate the performance of state-of-the-art deep generalized maxout networks, [16], in the context of multilingual and under-resourced monolingual speech recognition.

This paper is organized as follows: in section II we describe deep generalized maxout network training. Then, the phoneme mapping issues for multilingual DNN and both the knowledge-based and data-driven approaches are explained in section III. The databases and the experiments are presented in section IV and V. Finally we present concluding remarks.

#### II. DEEP GENERALIZED MAXOUT NETWORKS

A deep maxout neural network is simply a multilayer perceptron with many hidden layers before the softmax

<sup>\*</sup>This work is based on research supported by the South African National Research Foundation as well as the fund for scientific research of Flanders (FWO) under project AMODA GA122.10N.

<sup>&</sup>lt;sup>1</sup>Faculty of Electrical Engineering, KULeuven , 3001 Leuven, Belgium. Reza.Sahraeian@esat.kuleuven.be, Dirk.VanCompernolle@esat.kuleuven.be.

 $<sup>^2\</sup>mathrm{HLT}$  Research Group Meraka Institute, CSIR, South Africa. fdwet@csir.co.za

output layer and uses the maxout function to generate hidden activations [17]. Suppose  $\mathbf{u}^{(l)} = [u_1^{(l)}, u_2^{(l)}, ..., u_I^{(l)}]$  is a set of activations in layer *l*; where

$$\mu_i^{(l)} = \max_j (h_j^{(l)}), \ (i-1) \times g + 1 \le j \le i \times g$$
(1)

The function takes the maximum over groups of inputs,  $h_j^{(l)}$ s, which are arranged in groups of g.  $h_j^{(l)}$  is the *j*th element of  $\mathbf{h}^{(l)} = \mathbf{W}^{(l)}\mathbf{u}^{(l-1)} + \mathbf{b}^{(l)}$ .  $\mathbf{W}^{(l)}$  is the matrix of connection weights between the (l-1)th and *l*th layers,  $\mathbf{b}^{(l)}$  is the bias vector at the *l*th layer. In a maxout network, the nonlinearity is dimension-reducing and *I* is the dimensionality after the maxout function.

Generalized maxout networks may introduce other dimension reducing nonlinearities [16]. In this paper, we use the p-norm one:

$$u_i^{(l)} = (\sum_j |h_j^{(l)}|^p)^{\frac{1}{p}}, \ (i-1) \times g + 1 \le j \le i \times g$$
(2)

Where p is a configurable parameter. To train deep networks, greedy layer-wise supervised training [18] is used ; first, a randomly initialized network with one hidden layer is trained for a short time; then, the weights that go to the softmax layer are removed and a new hidden layer and two sets of randomly initialized weights are added. The neural network is trained again for the predefined number of iterations before a new hidden layer is inserted. This is repeated until we reach a desired number of layers. After the final iteration of training, the models from the last iterations are combined into a single model. In our study, the initial and final learning rates are specified by hand and equal to 0.02 and 0.004 respectively, and we always set p = 2. More details about the implementation and parameters are presented in [16].

# III. PHONEME MAPPING IN MULTILINGUAL DNN

Fig. 1 depicts the architecture of the typical multilingual DNN with shared hidden layers. In the multilingual target layer, each language can have its own output layer, Fig. 1-(a), or a common output layer is used Fig. 1-(b). In the latter, we need to provide a universal phoneme set; to this end, we may either consider a language label for each phoneme or merge phonemes. Simple concatenation of language specific phoneme sets, in the first scenario, may lead to performance degradation since very similar phones from different languages could be considered as different classes and the DNN would fail to discriminate between them [8]. For the second scenario, prior knowledge of a phonetician is required for the knowledge-based mapping which may not always be accurate and thus the DNN must encode disparate phonemes as a single class. This motivates us to investigate if a data-driven phoneme mapping can overcome the aforementioned problems. In the rest of this section, we describe the knowledge-based and data-driven phoneme mapping we used to train multilingual DNNs.

# A. Knowledge-based Phoneme Mapping

The major assumption for knowledge-based (KB) phoneme mapping is that the articulatory representations of



(a) Multilingual DNN with language dependent output layer.



Fig. 1. Multilingual DNNs with different types of output layers.

phonemes are similar and their acoustic realization can be assumed language independent. Based on this idea, universal phoneme inventories such as the IPA have been proposed [19]. In this study, the pronunciation dictionaries for the Afrikaans and Flemish include 37 and 47 phonemes respectively. In our KB phoneme mapping, each phoneme from the Flemish dictionary is mapped to only one of the phonemes in the Afrikaans one. To this end, 31 phonemes that share the same symbol in the IPA table are merged. However, there are 16 phonemes in Flemish without any IPA counterpart in Afrikaans which are mapped based on the linguistic knowledge. The phonemes:  $\tilde{\epsilon}$ ,  $\tilde{\alpha}$ ,  $\tilde{\sigma}$  and  $\tilde{\gamma}$  are simply mapped to / $\epsilon$  n/, / $\alpha$  n/, / $\sigma$  n/ and / $\gamma$  n/, and the rest are mapped as described in Table I.

#### B. Data-driven Phoneme Mapping

In our data-driven (DD) approach, we assume to have access to the pronunciation dictionary and the transcriptions for the target language. Then, each phoneme in Flemish can be mapped into *N*-best corresponding matches in the Afrikaans by calculating a confusion matrix.

Afterwards, a new pronunciation dictionary is created in which Flemish entries are described with the Afrikaans phonemes. Table II includes two examples explaining how

TABLE I Summary of knowledge-based phoneme mapping between Flemish(FL) and Afrikaans(Afr) languages.

Fl	Afr	Fl	Afr	Fl	Afr
Ŋ	х	YĽ	ə	ıc	Э
h	ĥ	OĽ	uə	εi	əi
Ι	3	e:	iə	13	3
Y	œ	aː	ar	au	əu

the Flemish words "met" and "stipt" are phonetized in the original Flemish lexicon and the new KB and DD ones. In the first example, the phoneme " $\epsilon$ " in the Flemish is mostly confused with three phonemes in the Afrikaans: "a", "ce" and "ai". Therefore, we consider three different pronunciations for this word based on the phoneme " $\epsilon$ " in the new lexicon. In this setup, the size of the new dictionaries increase rapidly with increasing N values. In addition, many of the Flemish phonemes have dominant matchings based on the confusion matrix; this is the case for almost all of the consonants. In this study, we set N=1 for the consonants and N=3 for the rest of the Flemish phonemes. It is also interesting to note that the Flemish phoneme "ɛ", for example, was merged with the Afrikaans phoneme of the same IPA symbol as in the KB phoneme mapping. However, "ɛ" is not among any of the three candidates chosen by DD approach. This indicates how differently the KB and the DD phoneme mapping may work.

In the second example, three different pronunciations for the word "stipt" are shown based on the phoneme "1". This phoneme has no IPA matching in Afrikaans and is mapped to " $\varepsilon$ " according to linguistic knowledge as shown in Table I. We should note that although the KB candidate for this phoneme is among those selected by DD approach, we have two more possible options for the mapping and depending on the context the best one will be chosen later based on the Viterbi alignment as a part of acoustic modeling. To

TABLE II New pronunciation modeling using DD and KB phoneme MAPPING.

Fl word	Fl lexicon	DD lexicin	KB lexicon
met(1)	mεt	m ə t	mεt
met(2)	-	mæt	-
met(3)	-	m əi t	-
stipt(1)	stipt	stept	stept
stipt(2)	-	stipt	-
stipt(3)	-	stəpt	-

generate the confusion matrix, we measure the KLD between distributions of phonemes:

$$D(P \parallel Q) = \int P(x) \log \frac{P(x)}{Q(x)} dx$$
(3)

Where P and Q represent density functions of the phonemes

distributions in Afrikaans and Flemish respectively. It is worth noting that since KLD is not symmetric, it is normally appropriate for P to be the reference distribution and Q to be an approximation to it [20]. KLD is straightforward for normal distributions. However, for the multivariate Gaussian Mixtures Models (GMMs), the KLD is not analytically tractable and therefore we can use the variational approximation of KLD between GMMs [21]:

$$D^{\nu}(P \parallel Q) = \sum_{a} w_{a} log \frac{\sum_{a'} w_{a'} e^{-D(P_{a} \parallel P_{a'})}}{\sum_{b} \hat{w}_{b} e^{-D(P_{a} \parallel Q_{b})}}$$
(4)

Where  $P = \sum_{a} P_{a}$  and  $P_{a} = w_{a}\mathcal{N}$ , and  $\mathcal{N}$  represents the normal distribution; similarly  $Q = \sum_{b} Q_{b}$  and  $Q_{b} = \hat{w}_{b}\mathcal{N}$ . w and  $\hat{w}$  are the Gaussian weights assigned to the Gaussian mixtures in the *P* and *Q* respectively.  $D^{v}$  is calculated for all pairs of phonemes in Afrikaans and Flemish to construct the confusion matrix. In this study, we use GMMs to model the phoneme distributions. Noting that the number of Gaussian components is set empirically and it equals two.

#### IV. DATABASES

# A. Afrikaans data

The NCHLT corpus<sup>1</sup> [22] is an Afrikaans database including broadband speech sampled at 16 kHz. The phoneme set contains 37 phonemes and silence. We have been provided with a pronunciation dictionary as well as training, test and validation sets. All repeated utterances were removed from the original dataset. In our setting, to simulate a low resource condition, a data set including 1 hour of data and 188 speakers was extracted from the training part and used together with the original validation and test sets (Table III).

TABLE III Description of the Afrikaans data set and a low resource subset for training purposes.

Set:	Train	Test	Dev
Duration	1h	2.2h	1.0h
# speakers	188	8	10

# B. Flemish Data

The Spoken Dutch Corpus (Corpus Gesproken Nederlands, CGN) is a standard Dutch database that includes speech data collected from adults in the Netherlands and Flanders [23]. This dataset consists of 13 components that correspond to different socio-situational settings. In this study, we used Flemish data (audio recordings of speakers in Flanders) from component-o which contains read speech. This dataset includes 38 hours of speech sampled at 16KHz and we have taken 36h for the training and 2h for the evaluation. In this work, we used only the training part including 36 hours as donor data produced by 150 speakers.

<sup>1</sup>Available from the South African Resource Management Agency (http://rma.nwu.ac.za/).

Flemish words in the CGN pronunciation dictionary are phonetized by 47 phonemes which are mapped to the 37 phonemes of Afrikaans.

# V. EXPERIMENTS

This section describes the experimental study performed to evaluate the impact of deep generalized maxout networks for low resource ASR as well as the proposed phoneme mappings for multilingual DNN training. First, monolingual experiments on Afrikaans are presented which serves as a baseline. Then, we used Flemish to improve this performance in the context of multilingual DNN. In this study, we used the Kaldi ASR toolkit [24] for DNN training.

# A. Monolingual Experiments

The first set of experiments was carried out on the Afrikaans language. We used a standard front-end by applying a Hamming window of 25ms length with a window overlap of 10ms. 13-dimensional features including 12 MFCC coefficients and the energy were extracted. Then, first and second derivatives were added and utterance-based mean and variance normalization was applied in both training and testing stages. These features were used to build 3-state left to right HMM triphone models with a total number of Gaussian components of  $\sim$ 3000; this value was set using the validation set (Table III).

We trained a bi-gram phoneme model on the training set and the ASR performance is reported in phoneme error rate (PER). The neural network's inputs were the 24-dimensional FBANK features being concatenated with 7 left and 7 right neighbor frames, yielding a 360 dimensional input layer; then, an LDA transformation matrix was applied without dimensionality reduction. We observed that FBANK features outperform MFCCs as input features for DNN. In this set of experiments, we first trained standard DNN systems with tanh activation functions. The number of context-dependent triphone states (i.e. DNN targets) is 505; the number of units in each layer equals 100 to achieve the best results. Table IV provides the ASR performance using both HMM/GMM and the corresponding hybrid DNN systems. Since we have only one hour of training data, increasing the number of hidden layers may degrade the performance. The PERs for hybrid DNN systems with 1 and 2 layers are reported in Table IV; we observed higher PERs for more hidden layers. The best performance for monolingual DNN with tanh nonlinearity is obtained with one hidden layer.

#### TABLE IV

PER(%) FOR AFRIKAANS USING HMM/GMM AND HYBRID DNN SYSTEMS WITH *tanh* ACTIVATION FUNCTION TRAINED ON AFRIKAANS DATA ONLY

	HMM/GMM	Hybrid DNN		
		1 layer	2 layers	
PER(%)	25.18	24.49	25.35	

Then, we trained DNNs with the p-norm activation function; in this case, we have one more parameter which is the group size, g. The proper value for g and other neural network parameters such as number of hidden layers and the input dimensionality for the p-norm activation are jointly tuned on the validation set. In Table V the PERs for different numbers of hidden layers and different values of g are presented. In these experiments I = 100 and various input dimensionalities are investigated. Table V shows that the performance is improved when a generalized maxout network is used for such a low resource setting.

#### TABLE V

PER(%) on the Afrikaans using hybrid DNN systems with p-norm nonlinearity and various settings where the p-norm output dimensionality is I = 400.

input dim	# of hidden layers				
input unit.	1	2	3	4	
400	23.61	23.83	23.68	23.72	
300	23.59	23.96	23.99	24.03	
200	23.76	23.71	24.01	24.01	

#### B. Multilingual Experiments

We subsequently merged the Flemish and Afrikaans training data based on both the knowledge-based and the datadriven universal phoneme sets explained in section III. Then, we trained a multilingual HMM/GMM system using 39dimensional MFCC features. The numbers of tied-states used for the multilingual HMM/GMM system are 4131 and 3973 for the KB and DD approaches respectively.

Table VI gives the performance of the multilingual HMM/GMM systems for the two types of phoneme mapping by using the same bi-gram language model trained with 1 hour of Afrikaans. These results are presented here to evaluate the effectiveness of the DD phoneme mapping. As shown, DD phoneme mapping considerably improves the performance of multilingual HMM/GMM systems; yet, it can be seen that the PER is much higher than the monolingual case presented in Table IV and Table V.

TABLE VI PER(%) COMPARISONS FOR KB AND DD PHONEME MAPPING USING A MULTILINGUAL HMM/GMM SYSTEM.

	KB mapping	DD mapping
PER(%)	45.89	39.81

Multilingual DNNs were subsequently trained by adopting context dependent decision trees and audio alignments from the multilingual HMM/GMM systems. In this set of experiments, the DNNs used *p*-norm activation functions and were trained from 15 consecutive frames and 24 FBANK features like DNN for monolingual setting. *p*-norm input and output dimensionality were empirically set to 1000 and 200 respectively. To bootstrap the acoustic model for Afrikaans, the hidden layers of the multilingual DNNs are shared and the softmax layer is replaced with the output layer corresponding to Afrikaans.



Fig. 2. PERs(%) comparisons for KB and DD phoneme mapping using multilingual DNN w.r.t. the number of hidden layers.

Fig. 2 shows a comparison of PERs obtained by multilingual DNNs with different numbers of hidden layers and reveals the following trends: first, both multilingual DNN systems provide significant reductions in ASR PERs when compared to the monolingual baseline systems presented in Table IV and Table V. Secondly, a comparison between the KB and DD phoneme mappings for DNN training shows that the ASR performance tends to improve in the case of using DD phoneme mapping. However, only marginal performance differences are observed if the neural networks are trained deep enough. This difference, however, depends on how similar the results of the two phoneme mapping techniques are. In this study, we observed that our DD technique maps all consonants to the same Afrikaans phonemes as the KB mapping does. Moreover, for many of the other Flemish phonemes, the selected KB candidate is among those chosen by the DD approach. For unrelated languages, however, DD phoneme mapping may perform differently and consequently lower PERs could be gained.

Finally, we examined another type of multilingual target where phoneme targets for Flemish and Afrikaans are kept separate Fig 1-(a). In this scenario, hidden layers are trained with data from both languages while the softmax layers are trained with language specific data where the number of output targets for Flemish is 4113 and 505 for Afrikaans.

TABLE VII PER(%) for 6 hidden layer multilingual DNNs with and without phoneme mapping.

	Phoneme mapping		No phoneme
	KB	DD	mapping
PER	18.29	18.25	21.04

Table VII shows that multileveled DNN approaches, either with or without phoneme mapping, improves ASR for lowresource languages. Moreover, we observe that phoneme mapping considerably improves the performance of multilingual DNNs. This can be due to the fact that Afrikaans and Flemish are closely related languages.

### VI. CONCLUSION

This paper presented an investigation of using generalized maxout networks and phoneme mappings for multilingual DNN based acoustic modeling. Our aim was to improve a speech recognizer for Afrikaans (as an example of a resource-scarce language) with generalized maxout networks and by borrowing data from Flemish (as an example of a related well-resourced language). Phoneme sets of these two languages were merged in both knowledge-based and data-driven fashions. We proposed to use an approximation of KLD to generate the confusion matrix for the DD phoneme mapping. This DD approach led to a performance improvement which was more pronounced in the multilingual HMM/GMM system than the DNN one. Moreover, we observed that if we train neural networks deep enough, the performance difference between two phoneme mapping approaches decreases. We also observed that phoneme mapping is beneficial when Flemish data is used to boost the Afrikaans recognizer in the framework of the multilingual DNN.

#### REFERENCES

- D. Imseng, P. Motlicek, H. Bourlard and P. Garner, Using outof-language data to improve an under-resourced speech recognizer, Speech Communication, vol. 56, 2014, pp. 142–151.
- [2] L. Burget, et al., Multilingual acoustic modeling for speech recognition based on subspace Gaussian mixture models, in Conf. Rec. 2010 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 4334–4337.
- [3] V. B. Le, and L. Besacier, First Steps in Fast Acoustic Modeling for a New Target Language: Application to Vietnamese, in Conf. Rec. 2005 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 821–824.
- [4] T. Schultz and A. Waibel, Language-independent and languageadaptive acoustic modeling for speech recognition, Speech Communication, vol. 35, 2001, pp. 31–51.
- [5] K. C. Sim and H. Li, Robust phone set mapping using decision tree clustering for cross-lingual phone recognition, in Conf. Rec. 2008 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 4309–4312.
- [6] W. Byrne, et al., Towards language independent acoustic modeling, in Conf. Rec. 2000 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. II1029–II1032.
- [7] J. T. Huang, J. Li,D. Yu, L. Deng and Y. Gong, Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers, in Conf. Rec. 2013 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 7304–7308.
- [8] K. Veselý, M. Karafiát, F. Grézl, M. Janda and E. Egorova, The language-independent bottleneck features, in Conf. Rec. 2012 IEEE Workshop on Spoken Language Technology (SLT), pp. 336–341.
- [9] S. Scanzio, P. Laface, L. Fissore, R. Gemello and F. Mana, On the use of a multilingual neural network front-end, in 2008 Proc. INTERSPEECH Conf., pp. 2711–2714.
- [10] N. T. Vu, et al., Multilingual deep neural network based acoustic modeling for rapid language adaptation, in Conf. Rec. 2014 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 7639–7643.
- [11] E. Egorova, K. Veselý, M. Karafiát, M. Janda and J. Cernocky, Manual and semi-automatic approaches to building a multilingual phoneme set, in Conf. Rec. 2013 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 7324–7328.
- [12] W. Heeringa, and F. De Wet, The origin of Afrikaans pronunciation: a comparison to west Germanic languages and Dutch dialects, in 2008 Proc. Pattern Recognition Association of South Africa Conf., pp. 159– 164.

- [13] F. Grezl, M. Karafiát and M. Janda, Study of probabilistic and bottle-neck features in multilingual environment, in Conf. Rec. 2011 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU), pp. 359–364.
- [14] P. Swietojanski, J. Li and J. T. Huang, Investigation of maxout networks for speech recognition, in Conf. Rec. 2014 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 7649–7653.
- [15] Y. Miao, F. Metze, and S. Rawat, Deep maxout networks for lowresource speech recognition, in Conf. Rec. 2013 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU), pp. 398– 403.
- [16] X. Zhang, J. Trmal, D. Povey and S. Khudanpur, Improving deep neural network acoustic models using generalized maxout networks, in Conf. Rec. 2014 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 215–219.
- [17] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville and Y. Bengio, Maxout networks, in 2013 Proc. ICML, pp. 1319–1327.
- [18] Y. Bengio, P. Lamblin, D. Popovici and H. Larochelle, Greedy layerwise training of deep networks, Advances in neural information processing systems, vol. 19, 2007, pp. 153–160.
- [19] International Phonetic Association, Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet, Cambridge University Press, 1999.
- [20] S. Kullback and R. A. Leibler, On information and sufficiency, The Annals of Mathematical Statistics, 1951, pp. 79–86.
- [21] J. R. Hershey and P. A. Olsen, Approximating the Kullback Leibler Divergence Between Gaussian Mixture Models, in Conf. Rec. 2007 IEEE Int. Conf. on Acoustics Speech and Signal Processing (ICASSP), pp. 317–320.
- [22] E. Barnard, M. H. Davel, C. van Heerden, F. de Wet and J. Badenhorst, The NCHLT speech corpus of the South African languages, in 2014 Proc. Workshop on Spoken Language Technologies for Underresourced Languages (SLTU), pp. 194–200.
- [23] N. Oostdijk, The Spoken Dutch Corpus. Overview and First Evaluation, in 2000 Proc. International Conference on Language Resources and Evaluation, pp. 887–894.
- [24] D. Povey, et al., The Kaldi speech recognition toolkit, in Conf. Rec. 2011 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU), pp. 1–4.

# Synthetic triphones from trajectory-based feature distributions

Jaco Badenhorst<sup>1,2</sup> and Marelie H. Davel<sup>1,3</sup>

<sup>1</sup>Human Language Technology Research Group, CSIR Meraka, South Africa. <sup>2</sup>Multilingual Speech Technologies, North-West University, Vanderbijlpark, South Africa. <sup>3</sup>CAIR, CSIR Meraka, South Africa.

jacbadenhorst@gmail.com, marelie.davel@gmail.com

Abstract—We experiment with a new method to create synthetic models of rare and unseen triphones in order to supplement limited automatic speech recognition (ASR) training data. A trajectory model is used to characterise seen transitions at the spectral level, and these models are then used to create features for unseen or rare triphones. We find that a fairly restricted model (piece-wise linear with three line segments per channel of a diphone transition) is able to represent training data quite accurately. We report on initial results when creating additional triphones for a single-speaker data set, finding small but significant gains, especially when adding additional samples of rare (rather than unseen) triphones.

**Index Terms**: synthetic triphones, trajectory modelling, trajectory-based features, feature distributions, feature construction

# I. INTRODUCTION

The accurate modelling of co-articulation effects in automatic speech recognition (ASR) systems has been a driving force behind the development of large speech corpora [1]. Whole word (or even phrasal) units capture co-articulation effects accurately within unit but require very large training corpora; limited training data forces the use of smaller units, and has resulted in the widespread use of context-dependent phones to capture co-articulation effects [1].

In practice, context sizes of three (triphones) or five (quinphones) are often used. When data is limited, many of these context-dependent units will rarely or ever be seen during training. In typical ASR systems, such unseen context-dependent units are modelled by clustering them with 'matching' seen units, based on a combination of acoustic and linguistic analysis, which is not always an optimal solution [2]. We are interested in determining whether it is possible to generate synthetic versions of such unseen or rare contexts from less specialised units observed in the training data.

First, we require a model that links more general units to more specialised units. For this purpose, we use a trajectory model that provides a compact way of representing the characteristic behaviour of transitions. From the characteristic trajectory behaviour of the less specialised transitions we reconstruct models for unseen transitions. In the current study, we restrict ourselves to triphone modelling, and aim to generate synthetic triphones from seen diphones. If this is possible, the same approach should be applicable to larger contexts, and possibly also to synthesizing additional speech data based on a small sample of data from a given speaker.

# II. BACKGROUND

In recent work, there has been renewed interest in data augmentation approaches for improving recognition accuracy for under-resourced languages. A useful way to group data augmentation schemes is to consider what type of additional data a technique produces. In [3], the three data types are referred to as 'other language', 'unsupervised' and 'synthesised' data. To incorporate other language data, systems utilise multilingual acoustic models trained on universal phone sets [4] exploiting resources across language barriers. Bootstrapping and filtering out the poor quality data is helpful (unsupervised techniques), while synthetic data may refer to perturbed data or entirely new examples of contexts which have been artificially generated. These techniques can potentially also generate vast amounts of data.

The work of both Jaitly and Hinton [5] and Kanda *et al.* [6] have shown that modelling accuracy can be improved by augmenting limited training data with synthetic training samples. For both cases, a modified version of the training data is added to the original data set when training hidden Markov models with deep neural networks. In the first case, vocal tract length normalisation (VTLN) is applied with different warping factors (the features are adjusted, labels kept unchanged) and in the second, different VTLN warping factors, different speech rates and frequency distortions are applied in a similar fashion.

Using trajectory models for the same goal, builds on prior work analysing co-articulation trajectories [7], [8], [9] as well as various studies on trajectory modelling for ASR purposes [10], [11], [12], [13]. Particularly, in [8] it was found that some decision-tree clustered triphones provided less accurate representations than a simple biphone model, providing the motivation for the current study.

As trajectory models, in effect, smooth features at frame level, they are related to the low-pass filtering used in noise robust speech recognition. The end goal of noise robust approaches is to systematically 'recover' corrupted speech frames. In principle, if the reliable features can be identified, these can then in turn be used to make more accurate predictions about less reliable ones. To this end, Chen and Bilmes [14] use Auto-Regressive Moving Average (ARMA) filtering at the cepstral level to improve ASR robustness in noisy conditions. If over-smoothing occurs, the more definite boundaries of speech events can be modelled using edge-



Fig. 1. Characteristic representation of a single transition (3-piece linear model).

preserved filtering [15]. Xiao and Li [16] show that besides normalising the probability distributions of speech features, the temporal characteristics of the feature trajectories can also be enhanced at the spectral level. In this work, we apply ARMA filtering at the spectral level, as a preprocessing step prior to fitting trajectories.

#### III. APPROACH

Generating synthetic triphones consists of four main steps: (1) fitting a trajectory model to seen transitions, (2) estimating trajectory parameters based on these seen transitions, (3) creating artificial utterances based on the estimated trajectory parameters, and (4) constructing trajectory-based features for model training.

#### A. Diphone segment-based trajectory model

We model the transitions that occur in speech with piecewise linear approximation at the spectral level. Three line pieces are used to fit a single feature channel of a filterbank, using least-squares optimisation. Figure 1 depicts this modelling strategy. A segment effectively describes a diphone, only using the closest 50% of monophone frames to the ASR boundary. We restrict the start and end line segments to be constant values (linear with zero slope), and model the transition between these two values with a straight line of variable slope. We require the constant line segments (the start and end line pieces, referred to as *stable values*) to be associated with at least 1 frame each. The connecting central line segment is referred to as the *change descriptor*.

Each stable value is estimated as the mean of associated feature values; the change descriptor is modelled by the first order line connecting the stable value anchor points. We optimise the squared error (SE) across all three line segments simultaneously by searching through the indexes of the possible start and end points for the change descriptor and draw the first order line between the end and starting indexes of the two anchor points. The squared errors at each instant are estimated, followed by the channel-specific mean

square error across frames:

$$MSE_{channel}(c) = \frac{1}{F} \sum_{f=1}^{F} |t_c(f) - x_{c,f}|^2$$
 (1)

where  $t_c(f)$  is the value of the trajectory function and  $x_{c,f}$  the true feature value, respectively, at frame f and feature channel c, and  $|t_c(f) - x_{c,f}|^2$  is the squared residual. F denotes the total number of frames for the segment. Once optimised, this model provides the following scalar values (see Figure 1):

- S1, S2 parameter value at initial and final stable value
- T1, T2 frame at start and end of the transition
- $T_{mid}$  centre of the transition
- $T_{dur}$  difference between T2 and T1 (2)

A similar SE measurement is also used to evaluate the extent to which trajectory models fit a set of speech data: the mean error (over all frames of all transitions) is now taken across all channels as well. Since channels have quite different standard deviations ( $\sigma_c$ ), a variance-weighted MSE ( $MSE_{weighted}$ ) is useful to evaluate:

$$MSE_{weighted} = \frac{1}{C} \sum_{c=1}^{C} \frac{1}{\sigma_c^2} MSE_{channel}(c) \quad (3)$$

with C the total number of feature channels.

# B. Predicting trajectory-based parameters

Given a set of training data, any group of segments can be modelled by a probability density function (pdf) over the parameters of section III-A. In this work, we choose to make the assumption that each parameter is normally distributed. The following section now describes how we use pdfs of the segment-based trajectory parameters to model speech data.

1) Estimating parameter distributions: Once an initial set of trajectories has been fitted to the training data, the mean and full-covariance matrix is estimated for the stable values (S1 and S2 respectively) of every particular biphone context that is required. Although biphone contexts are better resourced than triphone contexts, it is still not guaranteed that all biphones will have been seen. To supplement the under-estimated variances of these biphone values, we share the monophone diagonal variances for each distribution estimated on less than a fixed number of examples (3 in the current work).

Time alignments are modelled in a similar manner. Referring back to the schematic representation in Figure 1, these parameters are captured in two pdfs: (1)  $T_{dur}$  and (2)  $T_{mid}$ . A diphone context size is used.

2) Creating synthetic phones: As described above, we represent speech data at the spectral level with a set of four full-covariate pdfs (representing S1, S2,  $T_{mid}$  and  $T_{dur}$ ) per modelled context (currently diphones) and filterbank channel. For every triphone synthesised, the two diphones are created individually, by sampling from these four pdfs. Analytic



Fig. 2. Constructing a triphone model from two separate diphone transition segments.

constraints ( $seg_{start} < T1 < T2 < seg_{end}$ ) prevent invalid samples from being generated.

These synthetic diphones form the building blocks for generating artificial ASR data. Construction of a complete triphone example is now simply a matter of concatenating two appropriate diphone segments for a required context, as illustrated in Figure 2. Every triphone model requires shared stable values between the channel-based piece-wise linear models of two segments. We use the implementation described in [9] for this process, and form trajectories for complete utterances from the small diphone segments.

# C. Augmenting train data triphone classes

Additional triphone examples can now be added to the training data by generating fully artificial utterances. As it is not possible to simply stitch the required triphone labels together to form an utterance directly, we define an additional 'garbage' phone that can be added whenever subsequent triphone labels do not match. Each garbage model is generated from the preceding and following diphone segment pdfs, and is discarded after training, prior to decoding.

# D. Trajectory-based feature construction

In order to generate features from trajectory models, we extend the standard Mel-frequency cepstrum coefficient (MFCC) feature description: similar to generating standard MFCCs, the first step is to perform a fast Fourier transform (FFT) and obtain raw filterbank outputs. We use the Hidden Markov Model Toolkit (HTK) [17] and fairly standard parameters (sampling the speech signal at a frame rate of 5ms with a set of 26 filters).

Three additional steps are required before trajectories can be estimated: the log operation, mean subtraction for every channel and ARMA filtering [14]. Trajectories are only extracted for training data. (For the test data, we rely on ARMA filtering alone to smooth features. Alternatively, two pass-recognition can be used to obtain test trajectories and possibly further improve recognition accuracy; this is not evaluated as part of the current study.)

Sampling the trajectory models estimated at this point generates a new set of frame-based features with a standard 10ms frame rate. We apply the discrete cosine transform (DCT) with a cepstral liftering coefficient of 22 (as implemented in HTK). This provides 13 cepstral features, for which the standard first and second order derivatives are taken. Lastly, cepstral mean and variance normalisation (CMVN) are applied to the complete data set.

# IV. EXPERIMENTAL SETUP

All experiments use a single-speaker corpus, specifically designed for trajectory modelling: the Afrikaans Trajectory Tracking corpus (ATT) [9]. In the next section (Section V), we first show that the trajectory models approximate the training data fairly accurately. We then analyse the mismatch between triphones in the training and test sets and experiment with different ways of creating synthetic triphones, paying specific attention to the difference between reconstructed triphones (not seen at all in the training data) and rarely seen triphones.

# A. Speech data

The ATT corpus [9] consists of about 6 000 short utterances of a single male speaker, with a 4 974 subset considered of good audio and transcription quality. From this 'clean' data set, training and test data sets were selected of 4 072 and 902 utterances respectively, as described in [9]. As this data set is still quite large, we select a random subset of 961 utterances (less than 40 minutes of speech) to construct the under-resourced training data set. A further 440 utterances were selected (also randomly), as a development set.

# B. Segmentation and test system parameters

To segment the training data for trajectory modelling, we use a standard HMM-based ASR system trained on all 4 974 utterances and perform automatic alignment of the training data. (These alignments are not used during testing.)

In all experiments to follow, similar systems are trained: a context-dependent cross-word phone recogniser with tied triphone models and 39 cepstral trajectory features (the first 13 features as defined in Section III-D, and their first and second order derivatives). These features are computed with a window size of 25ms and at a frame rate of 10ms. Tied triphone models are estimated using standard phonetic decision-tree clustering. Each triphone model has 3 emitting states with 7 Gaussian mixtures per state and a diagonal covariance matrix; semitied transforms are applied. Only insertion penalties are optimised during decoding: for all experiments this is done using the development set, with results reported on the test set.

#### V. EXPERIMENTS AND RESULTS

A. Accurate trajectory-based representation

Nr	Feature type	Channels	MSE	$\rho$
1	Cepstral	13	0.1836	0.9133
2	Spectral	26	0.0633	0.9672
3	Spectral (ARMA 6)	26	0.0177	0.9834

TABLE I

Measuring approximation efficiency of trajectories with weighted MSE and correlation measurements.



Fig. 3. Number of triphone examples in train data (3754 triphone labels in total).

First, we evaluate the 'goodness of approximation' of the trajectory models by measuring the variance-weighted MSE between the original feature frames and the corresponding model values (using eq. 3). These values are shown in Table I for three different sets of features, evaluated on the test data set. We also calculate the Pearson correlation coefficient ( $\rho$ ) between trajectory estimates and actual feature values. From Table I it is clear that the two measures correlate well, that the ARMA-filtered spectral trajectories provide the best fit, and that a low MSE value (0.0177) and high correlation coefficient (0.9834) are obtained this way. The rest of the experiments only use ARMA-filtered spectral trajectories.

#### B. Triphone coverage

In Table II and Figure 3 the triphone coverage of the training data is given. 3 754 unique triphones are observed. On average, each is observed 5 times, with only about a  $\frac{1}{4}$  of labels occurring more than 5 times. Good overlap (2 608 labels) exists between the training and test data sets. Of the 1 045 labels not seen in the training set, 929 can be reconstructed from diphones (that is, both required diphones are observed in the training data).

Category	Triphone count
Train data	3 754
Test data	3 653
Test data seen	2 608
Test data unseen	1 045
Test data unseen: Diphone constructable	929
Test data unseen: Not constructable from diphones	116

TABLE	Π
-------	---

Triphone overlap between test and train data sets and that the number re-created from diphone transitions.

Figure 3 clearly shows how rarely seen triphones can also be exploited. We experiment with four pools of these labels: (1) < 2 examples with 692 labels, (2) < 3 examples - 1127 labels, (3) < 5 examples - 1616 and (4) < 10 examples where 2189 labels form part of the rare category.

# C. Baseline

Evaluating the phone recognition accuracy of the test data, leads to the baseline result ('Control') in Table III. Phone

accuracies are reported for systems with and without semitied transforms.

System	#Cons	#Rare	ACC	ACC (semitied)
Control	-	-	78.31	78.65
Reconstruct 1	1	-	78.66	79.54
Reconstruct 2	2	-	79.33	79.87
Reconstruct 3	3	-	78.76	79.40
Reconstruct 4	5	-	78.26	79.16
Reconstruct 5	10	-	76.54	79.08
Rare 1	-	2	79.26	79.70
Rare 2	-	3	79.95	80.04
Rare 3	-	5	79.56	80.39
Rare 4	-	10	77.34	79.31
Combined 1	2	2	78.47	80.55
Combined 2	1	3	79.48	80.82
Combined 3	2	3	79.46	79.99
Combined 4	3	3	79.39	80.40
Combined 5	3	5	78.89	80.30
Combined 6	5	5	78.53	79.54
Combined 7	10	10	75.96	78.33

# TABLE III

Phone recognition results when adding synthetic triphone examples to rare and unseen triphones in the training data.

#### D. Reconstructed triphones

We experiment with five acoustic models (Reconstruct 1 - 5), reconstructing unseen triphones; these only differ with regard to the number (#Cons) of reconstructed synthetic triphone examples generated per label (Table III). From these results, adding unseen triphones to the training data does improve phone recognition accuracy. Interestingly, only adding a few samples works well: adding too many examples (5 or 10) does not provide better accuracy.

# E. Rare triphones

Since at least one example of real training data is seen for the triphone identities in the rare triphone category, a smaller number of rarely seen triphone examples needs to be added to achieve the same set number of examples per triphone label. The results in Table III prove that the seen examples are not adequate: system accuracy improves significantly when adding synthetic triphone examples.

For the next four acoustic models, we steadily increase the number of synthetic triphones in the rare triphone category. In the Rare 1 experiment, at least 2 examples of all triphones are included (after additional triphones have been generated). Similarly, for systems Rare 2-4, at least 3, 5 and 10 examples are included in the training data. (See Table II for the number of triphones in each class.)

Again, adding 10 examples does not provide the best improvement (79.31%). The Rare 3 model provides a better result than that obtained for unseen triphones (Reconstruct 2 model).

#### F. Under-resourced triphones

Given the results above, the next question is whether generating synthetic triphone examples for both the rarely seen and unseen triphone categories would contribute to gains in phone recognition accuracy. In an attempt to do so, we combine the synthetic triphone example sets.

Forcing 10 triphone examples remains too many. The accuracy of 78.33% that the Combined 7 system achieves (Table III) show no gain over the baseline. Lowering the number of synthetic triphone examples to 5, significantly improves accuracy, but still does not outperform the previous Rare triphone experiments. In fact, the Combined 4 system results match these, but now for the system where we force a number of 3 examples of the under-resourced triphone label classes.

Since the reconstructed and rare triphone categories behave differently, as a last refinement, we also test what happens when different numbers of examples are added from each triphone class. We obtain the best results when a single triphone example is added for each of the 926 labels of the unseen category and at least 3 seen examples for the 1 127 labels of the rare category is forced. These results are tuned to the specific data set considered here: we do not propose it as a general strategy for adding synthetic triphones. Rather, we find it interesting that recognition accuracy can be improved using a fairly crude strategy for generating synthetic triphones.

# VI. DISCUSSION

When developing ASR systems in resource-constrained environments, many triphones are never seen during training. We found that it is possible to reconstruct unseen triphones from smaller contextual units, and that this can improve recognition accuracy of a standard tied-state baseline ASR system. Although speech synthesis techniques could also be used to generate unseen triphone units, we found that the trajectory models we use are able to represent training data surprisingly well. Our technique leads to a natural process for creating artificial utterances, containing repeated sequences of synthetically created samples of both unseen and rare triphones. Randomly selecting the training data set from the phonetically balanced ATT corpus [9], we can expect the distributions of unseen and rare triphones to remain comparable for new speakers of the same language.

With the current approach, the number of samples that can be added is still quite limited. Further refinements to the model are foreseen, especially with regard to the current sampling process, which is fairly crude. Trajectory optimisation and exploring the relationship of the new trajectories and the MSE or correlation measures in Table I may lead to improved results. Similarly, we report on improved results for semitied transforms, but the effect using training features with a high degree of feature smoothness on this technique remains to be investigated.

Our goal is to first analyse and understand the current restricted environment (a single speaker, generating triphones from diphones) in more depth, before considering the extent to which the current approach can generalise to larger contexts (quinphones from triphones) and finally, crossspeaker data augmentation. This is aligned with recent data augmentation approaches that have begun to address the cross-speaker problem. For example, [18] attempts to find a stochastic feature mapping (SFM) to statistically convert features between two speakers. The extent to which these approaches (generating additional synthetic contexts for a single speaker, and generating additional synthetic speakers) are complementary, remains an open question.

#### REFERENCES

- K.-F. Lee, "Large-vocabulary speaker-independent continuous speech recognition: The sphinx system," Ph.D. dissertation, Carnegie Mellon University, 1988.
- [2] H. Chang and J. Glass, "Multi-level context-dependent acoustic modeling for automatic speech recognition," in *Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on*, Waikoloa, HI, December 2011, pp. 89–94.
- [3] A. Ragni, K. Knill, S. Rath, and M. Gales, "Data augmentation for low resource languages," in *Proceedings of Interspeech*, Singapore, September 2014, pp. 810–814.
- [4] H. Lin, L. Deng, D. Yu, Y. Gong, A. Acero, and C.-H. Lee, "A study on multilingual acoustic modeling for large vocabulary ASR," in *Acoustics, Speech and Signal Processing (ICASSP), 2009 IEEE International Conference on*, Taipei, April 2009, pp. 4333–4336.
  [5] N. Jaitly and G. E. Hinton, "Vocal tract length perturbation (VTLP)
- [5] N. Jaitly and G. E. Hinton, "Vocal tract length perturbation (VTLP) improves speech recognition," in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
  [6] N. Kanda, R. Takeda, and Y. Obuchi, "Elastic spectral distortion
- [6] N. Kanda, R. Takeda, and Y. Obuchi, "Elastic spectral distortion for low resource speech recognition with deep neural networks," in *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on.* IEEE, 2013, pp. 309–314.
- [7] J. A. C. Badenhorst, M. H. Davel, and E. Barnard, "Analysing coarticulation using frame-based feature trajectories," in *Proceedings of PRASA*, Stellenbosch, South Africa, November 2010, pp. 13–18.
- [8] J. Badenhorst, M. Davel, and E. Barnard, "Trajectory behaviour at different phonemic context sizes," in *Proceedings of PRASA*, Vanderbijlpark, South Africa, November 2011, pp. 1–6.
- [9] J. Badenhorst, M. Davel, and E. Barnard, "Improved transition models for cepstral trajectories," in *Proceedings of PRASA*, Pretoria, South Africa, November 2012, pp. 157–164.
- [10] V. Digalakis, "Segment-based stochastic models of spectral dynamics for continuous speech recognition," Ph.D. dissertation, Boston University, 1992.
- [11] W. Holmes and M. J. Russell, "Probabilistic-trajectory segmental HMMs," *Computer Speech and Language*, vol. 13, no. 1, pp. 3–37, January 1999. [Online]. Available: http://www.idealibrary.com
- [12] L. Zhang and S. Renals, "Phone recognition analysis for trajectory HMM." in *Proc. Interspeech*, 2006.
- [13] D. Yu, L. Deng, and A. Acero, "A lattice search technique for a long-contextual-span hidden trajectory model of speech," *Speech Communication*, vol. 48, no. 9, pp. 1214–1226, September 2006.
- [14] C. P. Chen and J. A. Bilmes, "MVA processing of speech features," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 257–270, January 2007.
- [15] M. U. X. Lu, S. Matsuda and S. Nakamura, "Temporal contrast normalization and edge-preserved smoothing of temporal modulation structures of speech for robust speech recognition," *Speech Communication*, vol. 52, no. 1, pp. 1–11, January 2010.
  [16] S. C. X. Xiao and H. Li, "Normalization of the speech modulation
- [16] S. C. X. Xiao and H. Li, "Normalization of the speech modulation spectra for robust speech recognition," *IEEE Transactions on Computer Speech and Language*, vol. 16, no. 8, pp. 1662–1674, November 2008.
- [17] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, G. Moore, J. Odell, D. Ollason, D. Povey, V. Veltchev, and P. Woodland, *The HTK Book.* http://htk.eng.cam.ac.uk/: Cambridge University Engineering Department, 2005.
- [18] X. Cui, V. Goel, and B. Kingsbury, "Data augmentation for deep neural network acoustic modeling," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference.* IEEE, 2014, pp. 5582–5586.

# Review of Standard Rotor Configurations for a Micro Unmanned Aerial System

Angus Steele<sup>1</sup> and Johann Treurnicht<sup>2</sup>

Abstract— The use of unmanned aerial systems (UAS) is on the rise with an array of industries finding use for them in a variety of applications. This review hopes to assist potential drone designers in selecting the drone best suited for their application. This paper attempts to first give a better understanding of flight theory and the basics of rotary winged vehicles. Next it builds on that knowledge and applies it to a few important selection parameters, after which it addresses the criteria and links it directly to a few standard configurations of rotorcraft. In the final discussion a few key points are addressed and each standard configuration is discussed in its ideal application.

#### I. INTRODUCTION

With the increase in computing power making rotorcraft control possible, drones are becoming an ever more popular platform for a variety of applications and users. From military personnel to photographers, drones are benefiting an array of industries. The hobbyist now can pick and choose from a warehouse full of different rotors, motors and full drone kits.

There is a lot of documentation available about rotorcraft, especially since they has recently gained a lot of attention from hobbyists. Unfortunately there are not a lot of comparisons between the different types of configurations, nor what parameters should be looked at. This misunderstanding results in many users end up using a drone type that is not the optimal choice for the application. This paper will summarise a few key points of rotor dynamic theory and then use this information to address a few important parameters before finally applying it to different drone configurations. By the end of this paper the reader should be able to identify the appropriate rotor configuration for their application.

# II. FUNDAMENTALS OF FLIGHT THEORY

# A. Basic Rotor Theory

The rotor is responsible for all the aspects of flight and generates the lift, forward propulsion and the means to control the orientation of the craft [15]. It is for this reason that an in depth understanding of rotor characteristics and performance is needed. Any rotating blade will cause a rotation of the craft in the opposite direction to that motion. This applied moment must be countered by a counter-torque mechanism. This function is performed by the tail rotor in a traditional helicopter.



Fig. 1. Velocity components of a rotor (taken from [15])

The capability of any part of a rotor to produce lift is influenced by the local blade position and pressure at that point [15]. As the rotor spins, the blade's angle relative to air stream in shifts as shown Figure 1. This angle is defined as an azimuth angle ( $\alpha$ ) and is measured relative to air flow. The azimuth angle is 0° downstream and sits at 180° when it faces directly upstream. The speed of any part of the rotor varies along the length of the rotor, with the maximum velocity at the rotor tip. Figure 1 demonstrates the naming scheme.

As the rotorcraft adds a horizontal component to its hover or vertical flight, the relative speed of the individual rotor segments now adheres to (1). The relative velocity at any part of the rotor is affected by the azimuth angle of the blade ( $\alpha$ ), forward translatory speed of the craft ( $V_{\infty}$ ), angular speed of the rotor ( $\Omega$ ) and the considered distance along the rotor blade (r) [15], [1].

$$V_r = \Omega r + V_\infty \sin(\alpha) \tag{1}$$

What this relationship shows is that during forward flight the tip velocity, relative to the ground, changes even if the rotor rotates at a constant speed. This relationship complicates the rotor dynamics at higher speeds and limits the top speed of the craft. On the retreating edge ( $\alpha = 270^{\circ} \therefore \sin(\alpha) = -1$ ) if  $\Omega r \leq V_{\infty}$  the rotor would effectively be going backwards and the helicopter is at risk of stalling out. This relative speed is known as a stall condition, while the advancing edge is reaching its maximum speed by approaching transonic conditions and severe instability [15], [1].

#### B. Momentum Theory and the Basics of Thrust

As mentioned above, the rotors of a rotorcraft are responsible for generating all the forces that manoeuvre the

<sup>\*</sup>This work was supported by the CSIR

<sup>&</sup>lt;sup>1</sup>Angus Steele is an employee at the CSIR as well as an MSc Student at The University of Stellenbosch (ASteele@csir.co.za)

<sup>&</sup>lt;sup>2</sup>Johann Treurnicht is with the Department of Electronic Engineering, The University of Stellenbosch



Fig. 2. Momentum Theory in Hover (adapted from [15])

vehicle. These forces are induced by pushing air through the rotor disk. With a fixed wing aircraft the analysis of the blades is simplified because the only air flow produced is from the translational velocity of the entire craft. Analysis of blade performance in a rotorcraft can be more challenging as the rotation of the blades must be considered alongside the overall speed of the vehicle. As the craft manoeuvres in space, the air flow through the rotor adds significant complexities to the analysis. Since the rotorcraft is expected to perform in a variety of flight styles it is important to understand these models as well as their flaws.

Thrust is the name given to the collective vector with lift being the component that opposes the weight, for simplification the helicopter is considered to be in a state of hover (weight = lift = thrust). The rotor 'smooths' out the air as it forces it through the disk area. This more uniform air creates an edge known as the slipstream or wake boundary, with the surrounding air remaining dormant [15]. Inside the wake boundary, the average velocity of the air is tangible, where outside the slipstream edge, the average air velocity is negligible. The force required to push that mass of air through the disk space is, by Newton's third law, returned by the air to the rotor, thus giving the rotor blades a thrust component.

Rankine-Froude's Momentum Theory looks at this induced velocity as well as the displacement of air through the propeller, and attempts to quantify the induced thrust. The variable naming convention for the equations is shown in Figure 2 above and is adopted from Leishman et al [15] in their naming of the components. Subscripts 0, 1, 2 and  $\infty$  refer to the locations of quiescent flow, inflow directly before the rotor, airflow immediately after the disk and the slipstream<sup>1</sup> or far wake condition respectively. The velocities are shown as the induced velocity in and out the rotor ( $v_i$ ), the far wake velocity ( $v_{\infty}$ ) and finally  $v_0$  represents the zone with zero flow rate. There is no velocity discontinuity across the rotor, the energy being fed into the system by the rotor is represented by a pressure change between  $P_1$  and  $P_2$ .

As described above, it is by forcing the air through the

<sup>1</sup>Generally far wake is considered as 1 full rotor diameter distance away [15].

disk that lift is generated. The mass flow rate of this air can then be described by (2), where  $\rho$  is the density of air and A is the area of one full blade rotation. The rate at which this mass of air is displaced becomes a crucial variable in rotor dynamics and is directly proportional to thrust (T). This relationship can be quantified as shown in (3). Thrust can also be calculated by finding the difference in pressures over the rotor disk as in (4)

$$\dot{m} = \rho A v_i$$
 (2)

$$T = \dot{m}a$$
 (3)

$$T = A(P_2 - P_1) \tag{4}$$

Since  $v_0$  is zero during hover and acceleration is the difference in  $v_{\infty}$  and  $v_0$ , (5) can be obtained.

$$T = \rho A v_i v_{\infty} \tag{5}$$

Thrust can now be quantified once the slipstream and induced velocities are known. Then by applying Bernoulli's equation of conservation to both sides of the rotor disk, the change in pressure across the disk can be quantified as shown in (6). That change in pressure fits into one of the initial definitions of thrust (4). Equating both of those definitions yields an important relationship between the three velocities, as shown in (7). The relationship simply states that the induced velocity at the rotor is the average of the quiescent flow above and the far wake velocities. This definition proves useful at a later stage in the rotor theory definitions.

$$P_2 - P_1 = \frac{1}{2}\rho(v_{\infty}^2 - v_0^2)$$
(6)

$$v_i = \frac{1}{2}(v_{\infty} + v_0)$$
 (7)

# C. Disk and Power Loading

Disk loading (DL) is a term seen often in the world of rotorcraft, and describes a simple but important ratio between thrust and the area a rotating disk covers. It is represented in its simplest form in the beginning of (8) and is measured in  $\frac{N}{m^2}$ . Since the pressure drop across each rotor is considered uniform, the disk loading for each rotor will equate to the pressure drop across that disk. Equation (6) first shows the difference in pressure and by taking  $v_0$  as zero (state of hover), the second half of equation (8) can be formed.

$$DL = \frac{T}{A} = \frac{1}{2}\rho v_{\infty}^2 \tag{8}$$

For multi-rotor craft, the disk loading is assumed uniform across all rotors [15]. The overall disk loading of a single rotor craft such as a traditional helicopter will be lower than that of a multi-rotor craft of a similar size [1]. Figure 3 shows some examples of disk loading values for a variety of rotor configurations, as shown, disk loading is also a measure of hover efficiency. A higher disk loading value results in larger values for induced velocities as well as the required power to hover. This relationship means that the larger the blades, the better the efficiency. More force will be generated by pushing large quantities of air slowly than forcing small amounts of air through at high speeds [21]. With bigger blades comes larger rotational inertia and geometry. The craft is also less immune to gusts and interferences. A larger blade creates faster tip velocities which will limit the speed of the craft severely [15].

Power is given by the product of thrust and the induced velocity at the blade. It can be written as shown in (9). What this ratio shows is that the ideal power is in cubic proportion to the induced velocity at the rotor. Therefore to reduce required power the rotors induced velocity must be small, which can be accomplished by an increase in disk area [15].

$$P = 2\rho A v_i^3 \tag{9}$$

Another important ratio is between thrust and power. It is called power loading (PL) and is shown in (10). Power loading can be seen as a measure of craft efficiency and is measured in  $\frac{N}{LW}$ .

$$PL = \frac{T}{P} \tag{10}$$

From (8) and (10) it can be shown that power loading is inversely proportional to disk loading. Therefore a craft with lower disk loading will be a more efficient platform.

# D. Electrical Power to Thrust

Equation (9) gives a quantitative approach to solving for aerodynamic power ( $P_i$ ). If electrical power is taken as  $P_e = VI$ , where V is the applied voltage and I is the sourced current, with an efficiency of  $\eta$  then  $P_i = \eta VI$ . Noting that  $P_i = Tv_i$  and using (9), a relationship between thrust and  $P_e$ can be formed and is represented in (11).

$$T = (2\rho A)^{\frac{1}{3}} (\eta P_e)^{\frac{2}{3}}$$
(11)

Equation (11) brings to light a very important relationship which states that thrust grows at a slower rate than the electrical input power to the system.



Fig. 3. Image representing, various Disk Loading values for varying rotorcraft (Taken from [15])

# $T \propto P_e^{\frac{2}{3}}$

# **III. SELECTION PARAMETERS**

Some of the fundamental theories described above relate to the basics behind various rotor configurations and even varying flight techniques. Each different arrangement of blades introduces certain advantages and disadvantages to the system. Not every configuration will be applicable for all operations and it is important to determine what criteria are critical for the intended application. An analysis of varying rotor configurations is done below and follows a similar trend to that seen in [5], [2] and [22]. The main weighted criteria for the discussion were listed in no particular order as:

- 1) Flight time, payload capability and efficiency
- 2) Geometry, size and mechanical complexity
- 3) Drone manoeuvrability and control algorithms
- 4) Stability and disturbance rejection

#### A. Efficiency, Payload Capability and Flight Time

Flight time is a by-product of efficiency and payload, as a more efficient craft will drain the power source slower, thus producing a longer flight. A heavier craft, with a larger payload, will have a shorter flight duration. Most miniature drones today have a flight time between 6 - 12 minutes before they need recharging. This length is not suitable for a variety of applications where longer mission durations are pertinent. A larger power source could always be added, but will increase the weight, limiting the payload capability and once again flight time. The relation between hover efficiency and disk loading was mentioned above, but what was not discussed is how a potential payload affects these decisions. With no payload a single rotor will have a lower disk loading and will not be able to carry as heavy loads as a multi-rotor craft that has a higher disk loading [15], [13], [1].

#### B. Geometry, Size and Mechanical Complexity

In any aerial vehicle mass is always an important design criterion. Every aspect of the platform must be designed to be the lightest it possibly can. Having a light weight craft is one part of the design criterion, another would be ensuring that the weight is geometrically spread out correctly, as well as functionally distributed appropriately. The table below was adapted from [22] and demonstrates the latter point. Depending on the different criteria for the craft, different functional blocks will be allocated a certain percentage of weight. For example, if the user would like a longer flight time, a higher percentage would be given to the power source and possibly less to the external payload. Generating a good mass model before designing helps better understand the requirements for the craft and could be a deciding factor in the construction.

Component	0.3kg	1.8kg	3.7kg
Rotor System	11.0	11.2	13.9
Tailboom Assembly	8.0	9.1	7.8
Main Rotor Motor	15.4	10.5	8.1
Fuselage/Structure	7.0	15.1	12.0
Main Transmission	2.0	3.4	3.4
Landing Gear	2.3	3.4	2.9
Control System	5.7	18.3	9.3
Avionics	29.4	2.4	1.6
Power Source	19.2	26.6	41.0

TABLE I

UAS WEIGHT DATA (ADAPTED FROM [22])

It was also mentioned that the weight needs to be geometrically positioned correctly, the point of this would be to create as much symmetry in the craft as possible. If this is done correctly the principal axes of inertia will align very closely with the body of the craft. The inertia tensor is a matrix that is a representation of a rigid body's resistance to movements in three dimensional (3D) space. For the general case the inertia tensor takes the form as shown in (12). The inertia tensor is very dependant on a craft's symmetry, and is symmetric itself. In other words,  $I_{xy} = I_{yx}$ ,  $I_{xz} = I_{zx}$  and  $I_{zy} = I_{yz}$  and therefore if a craft is symmetric about the y axis (x = 0), then  $I_{xy} = I_{yx} = 0 = I_{xz} = I_{zx}$  [17], [6].

$$\mathbf{I} = \begin{bmatrix} I_{xx} & -I_{xy} & -I_{xz} \\ -I_{yx} & I_{yy} & -I_{yz} \\ -I_{zx} & -I_{zy} & I_{zz} \end{bmatrix}$$
(12)

Symmetry in a craft can also help reduce the effects of disturbances such as mechanical drag and even wind. More on disturbance rejection is mentioned below.

#### C. Drone Manoeuvrability and Control Algorithms

In 3D space there are effectively six possible degrees of freedom (DOF) three of which are translational ( $\xi$ ) and three of them rotational ( $\eta$ ). The naming scheme used in this paper follows the same form used by Castillo et al in [6], [7]. In mathematics there are discussions regarding rotations around the x, y and z axes, though in flight theory they are labelled as roll, pitch and yaw. The three axes change as the aircraft's orientation changes since they are labelled relative to the aircraft's position. Pitch relates to how much the vehicle is tipping forward or backward, roll is an influence in the left and right rotation, while yaw is rotation around the z axis. Instead of x, y and z, these axes can be considered as forward, sideways and normal [15]. Refer to Figure 4 for a visual description of the axes.

$$\boldsymbol{\xi} = \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \\ \boldsymbol{z} \end{bmatrix} \quad \boldsymbol{\eta} = \begin{bmatrix} \boldsymbol{\phi} \\ \boldsymbol{\theta} \\ \boldsymbol{\psi} \end{bmatrix} \quad \boldsymbol{q} = \begin{bmatrix} \boldsymbol{\xi} \\ \boldsymbol{\eta} \end{bmatrix}$$
(13)

Drone manoeuvrability and control algorithms have been grouped together because they both relate directly to the dynamic model of the craft, as well as the amount of control authority available to the pilot. The same way that the wheels in a car decide which direction the car drives, the rotors provide all the control authority to a standard rotorcraft. Having only a single, fixed pitched rotor allows only for control in the amount of upwards and downwards acceleration the craft has. There are many different methods to obtain full six degrees of flight freedom.

Typically a rotorcraft will be designed with either fixed pitched, or variable pitched rotors. A fixed pitched rotor is a rotor that has an optimally selected, unchangeable pitch and therefore a fixed angle of attack. This means that since the angle of attack is fixed for the blade, an increase in RPM will be required for a change in lift. With a variable pitched blade, the pilot can change the angle of attack to increase the forces. As the angle of attack increases, the blade will produce more lift without changing the speed of the motor. However, as the pitch increases, so does the drag of the blade. This drag then requires more motor power to keep the blade moving through the air [6], [15].

The power requirements for either system are fairly similar; though the advantages of a varying pitch is that a single rotor has the potential for more dynamic force applications. The downfall however is the high level of complexity in the mechanical design. Both of these facts become pertinent in the final platform design. The end goal is to have a craft that can fly stably and accurately in three dimensions. To do this the craft will need more control surfaces to apply forces in those planes. There are many different methods to obtain the full six degrees of flight freedom. Some designers have added multiple rotors, ensuring there is always a counter rotating pair which eliminates the anti-torque generated by each motor. Ultimately giving the engineer more control authority will simplify the control algorithms and increase drone manoeuvrability. Having only a single, fixed pitched rotor will allow only for control in the amount the craft flies up or down, as well some instability in the system.

Most configurations will give the user sufficient control authority, the trade off becomes between number of rotors and mechanical complexity.

# D. Stability and Disturbance Rejection

Stability and disturbance rejection are generally considered control problems and a good control system can stabilise



Fig. 4. Control surfaces required for 3 dimensional flight (Taken from [10])

the craft amongst disturbances. These parameters have been isolated in this case to focus on what can be done before there is an attempt to apply control theory. Stability is a broad term and what is meant by it in this case is the ability to completely control movements in all six DOF. Any rotating member will produce a counter rotating torque to the static body, which means that any system with only one fixed pitched rotor will have inherent instability in the yaw axis and only vertical control [6]. It was mentioned earlier that symmetry in the craft can help eliminate the effects of some disturbances. Multiple blades helps reduce the effects of disturbances just as well. This way if one rotor falters or is affected substantially by a disturbance the other rotors can still rectify the error. Some multi-rotor designs can still fly with substantial control even after losing power to one or more of the rotors [18].

# **IV. DISCUSSION**

This discussion looks at the different rotor configurations and attempts to address each parameter mentioned above. It begins with the traditional helicopter which is always seen as a main rotor with a smaller rotor at the tail, even when there are many different types of anti torque tail set-ups. The ducted fan approach reduces the risk of damaging the tail rotor while silencing the system. The NOTAR design [16] manipulates the airflow generated by the main rotor and directs it to counter act the induced torque. A tip-jet design eliminates the torque applied to the airframe and therefore no tail rotor is required [5]. There have been many attempts at improving the standard helicopter design. These improvements have taken the form of adding rotors, designing hybrid aircraft and complex mechanical designs to harvest advantages of both the fixed wing and vertical take off and landing (VTOL) craft. Some have even tried to combine multiple features as Flanigan [11] did in his design of a tip-jet, compound, tilt rotor aircraft.

In an attempt at simplification, not all configurations were investigated. The following standard groups of designs were covered:

- 1) Traditional helicopter
- 2) Coaxial rotors
- 3) Tandem rotors
- 4) Multirotor designs
- 5) Tilt rotors

#### A. Traditional Helicopter

When most people think of a rotorcraft they will think of a conventional helicopter, which is still the most widely used configuration for large rotorcraft [5]. It consists of a single main rotor, coupled with a smaller rotor located in the tail to counter act the developed counter torque.

The main rotor of a standard helicopter has very low disk loading which gives it excellent hover efficiency. To achieve yaw stability this configuration makes use of a small tail rotor to counter act the induced moments of the main rotor. The extended tail rotor requires energy which it will draw from the motor while also adding a significant amount of length to the craft. Since the single rotor only gives the pilot thrust control and the tail rotor gives measurable yaw control, there is need for more control surfaces to do more manoeuvring. Most helicopters use a variable pitched rotor system. Cyclic control of this pitch allows the pilot to adjust the angle of attack of the rotor blades while they rotate. This set up is mechanically very complex but luckily has become a standard production set up, with many companies providing solutions to this problem.

Once the mechanics are set up the control algorithms are still slightly limited and intensive. However with only a single main rotor the traditional helicopter is extremely susceptible to disturbances and has a limited payload capability with the low disk loading factor. The need for the tail-boom assembly also adds significant length to the craft.

# B. Coaxial Rotors

A coaxial configuration consists of two counter rotating blades located about the same centre of rotation that both use the same drive system. This counter rotating pair eliminates the need for a tail rotor as the torque applied by both rotors cancel. Functionally the coaxial is very similar to the traditional helicopter [8]. With no modifications and only using fixed pitched rotors, this platform will only give yaw and over all thrust control. Bohorquez et al [2] attempted a number of lateral control methods, eventually settling on aerodynamic flaps to control the flow of the downwash. That, and other methods, are shown in Figure 5. Briod et all also used the same set up in his team's design of the Gimball [4].

The control flaps are the most common form of lateral control for a small coaxial UAS. They introduce little mechanical complexity and do not require excessive power to use. The flaps do however decrease efficiency of the system, but if designed correctly should only influence the system while being used. For hover and vertical flight the impact will be negligible. As a control surface the flap is quite rudimentary and will require more advanced control methods as well as in depth testing to obtain smooth flight transitions. Due to its compactness the design can have considerable manoeuvrability if the control algorithms are designed effectively. Each flap will require an actuator, which will increase total weight, power consumption and required mechanics.

Since the bottom rotor is working in the top blades slipstream, it will have a higher  $v_0$  and therefore a larger  $v_i$ , which according to equation 3 will induce a larger thrust. This coupling relates to high values of efficiency and lower values of disk loading, decreasing the payload capability.



Fig. 5. Different methods of lateral control in a Coaxial UAS (Adapted from [2])

Coleman in [8] did an extensive survey of coaxial rotors and also found that they produce more drag than the conventional rotor set up, which becomes pertinent at higher speeds.

Localising the blades around a single point also helps with the geometry of the craft as it is a more compact design. Briod et al [4], [14], [3] used this to their advantage when they were designing a collision resistant robot. The compact design allowed them to surround the entire craft in a rolling protective cage. The set-up of the coaxial main rotors do make the craft vulnerable to disturbances [8].

#### C. Tandem Rotors

A tandem rotorcraft is sometimes referred to as a dual rotor, as it consists of two blades to generate lift and decrease disk loading while increasing the payload lift capacity. In a tandem configuration the blades sit in the front and the rear of the craft, sometimes slightly overlapping. Tandems are often used in applications that require heavier loads than the traditional helicopter can effectively offer. In a tandem configuration the blades spin in opposite directions to counteract the other one's rotational torque. To obtain control in all wanted degrees of freedom the tandem requires the use of variable pitch rotors, similar to the system used in a traditional helicopter. In the case where the rotors overlap, a tandem helicopter has the problem of its rear rotor being influenced by the wake of the front rotor [13]. Running both rotors at a constant speed limits the chance of fatal rotor collision.

As described in (11) the thrust of the system increases slower than the electrical power input into the system. In a standard configuration, doubling the electrical power would only increase the thrust by a factor of  $\approx 1.587$ . Whereas doubling the amount of rotors being driven will double both the thrust and the electrical power. This relationship gives the tandem arrangement the capability of lifting heavier loads with relatively low power consumption, as well as demonstrating low power consumption for hover and slow translatory flight. Having twin blades does increase the size of the craft, but the elimination of the tail rotor sees the size being similar to that of a classic helicopter. Using two blades also decreases the effects of interferences such as gusts on the craft.

Coupling of the two gear boxes helps prevent inter-rotor collisions but that does not negate the need for two rotors requiring cyclic control, creating a system that requires complex control.

#### D. Multirotor Designs

Drones have joined other remote controlled vehicles in the world of hobbyists. Of all the different designs, the multirotor is the most popular. Through discussions with drone designers and aerial photographers, the four rotor design is generally chosen due to its incredible stability, manoeuvrability and disturbance rejection.

A quadrotor consists of two pairs of counter rotating propellers. Each shaft will be driven by its own motor and unlike the flaps in a coaxial system, every motor in



Fig. 6. Quadrotor vs Helicopter Rotor Spacing

a quadcopter attributes to the lift vector. Having the freedom to control each blade independently gives the pilot advanced manoeuvrability, with minimal mechanical complexities. This configuration also reduces the complexity of the control algorithms as four degrees of freedom can be obtained by simply adjusting the speed of the motors, with two secondary DOF. The multirotor can even rotate on the spot without any change in altitude [19].

The quad does however have very high disk loading due its utilisation of rotor space, Figure 6 demonstrates the concept. Imagine a traditional helicopter with radius R = 100 which creates an area  $A = \pi \times 100^2$ , to fit 4 rotors in the same area, without over lap would require a rotor radius of  $r < \frac{R}{2} \approx 40$  relating to a total area of  $4\pi \times 40^2$ , reducing the total disk area by about 1.5. This reduced disk area with the increase in weight relates to a higher disk loading and a less efficient hover; thus a more power hungry system. Similar to the tandem the increase in number of rotors increases the payload capacity of the multirotor and allows it to lift heavy loads. There are even products that have eight rotors to seriously increase the payload capability [9].

Besides the poor hover efficiency, the biggest downside of the multirotor designs is their size and weight. Each blade requires a drive system and space to rotate without interference.

# E. Tilt Rotors

A tilt rotor is a very sophisticated system that attempts to harness the benefits of both the fixed and rotor wing aircraft. With the addition of a pivoting axis for each blade the craft has the forward flying speeds of a fixed wing craft while still being able to take off and land vertically like a rotorcraft. The tilt rotor's major downfall is related to the required highly complex and intricate mechanical design [5].

VTOL applications require a larger blade to decrease the disk loading, while in forward flight a smaller diameter blade is desired to increase the efficiency of propulsion. Hager [12] developed a telescopic system that transforms the blades to get the optimal benefits out of each configuration. These and other improvements have established the tilt rotor as a competitive design in the field of aeronautic transportation [5]. The main advantage of the VTOL system compared to other rotorcraft is the flight efficiency in longer flights.

# V. CONCLUSION

From the above discussion it can be seen that there is no single configuration that is best suited for all applications. Instead each set up has its own pros and cons which can be utilised in different situations. To conclude this review, each standard configuration will be discussed in its ideal application.

Starting with the most application specific, the tilt rotor will only be the most advantageous in a situation that requires flight duration over long distances. With the added need of VTOL, the tilt rotor will trump the conventional fixed wing design. As expected the tilt rotor is also the best choice when it comes to top lateral speed.

The traditional helicopter and the coaxial fulfil a very similar role. They can be sensitive to disturbances and can't handle the payload their multirotor relatives can. However, their hover efficiency is very high which gives them a significant flight time and that's where the traditional setup earns its place. When an application's main criteria is extended flight time, these should definitely be considered. The choice between coaxial and traditional comes mainly down to size and flight speed. The coaxial will be able to fit in more refined spaced without the additional tail boom assembly, while the traditional will be able to reach higher speeds and fly laterally more efficiently.

The multirotor is the easiest to use, can take the biggest payload and is the most stable, but it will have a shorter flight duration as it is a power hungry system. For the case of the aerial photography it is no surprise that the multirotor was chosen as the best choice as flight duration is not as important to a photographer as stability would be. The multirotor is also the easiest to control which makes it the ideal hobbyist platform.

The tandem has often been cast aside as a suitable configuration [6], mainly because it sits between the traditional and the multirotor on effectively every parameter. So generally one or the other is chosen and the dual rotor set up is neglected. The tandem is suitable for an application that needs a jack of all trades solution. It is a slightly simpler system and will provide a larger payload than the traditional helicopter, while it still has better hover efficiency and overall size compared to the multi-rotor.

# ACKNOWLEDGMENT

The authors of this paper would like to extend a thank you to the Council of Scientific and Industrial Research (CSIR) for allowing the space and resources to conduct this study, as well as The University of Stellenbosch for all the advice and direction.

#### REFERENCES

- F.A. Association. *Rotorcraft Flying Handbook*. FAA Handbooks Series. Aviation Supplies & Academics, Incorporated, 2001.
- [2] Felipe Bohorquez, Paul Samuel, Jayant Sirohi, Darryll Pines, Lael Rudd, and Ron Perel. Design, Analysis and Hover Performance of a Rotary Wing Micro Air Vehicle. *Journal of the American Helicopter Society*, 48(2):80, 2003.

- [3] Adrien Briod, Dario Floreano, Przemysław Kornatowski, and Jean-Chirstophe Zuffery. A Collision-resilient Flying Robot. *Journal of Field Robotics*, 31(4):496–509, 2014.
- [4] Adrien Briod, Adam Klaptocz, Jean-christophe Zufferey, and Dario Floreano. The AirBurr : A Flying Robot That Can Exploit Collisions. pages 569–574, 2012.
- [5] Yihua Cao, Dong Li, Qiang Zhang, and Hang Bian. Recent Development of Rotorcraft Configuration. *Recent Patents on Engineering*, 1(1):49–70, 2007.
- [6] Lozano R Dzul AE Castillo, P. Modelling and control of mini-flying machines.
- [7] Pedro Castillo, Alejandro Dzul, and Rogelio Lozano. Real-time stabilization and tracking of a four-rotor mini rotorcraft. *IEEE Transactions on Control Systems Technology*, 12(4):510–516, 2004.
- [8] Colin P. Coleman. A Survey of Theoretical and Experimental Coaxial Rotor Aerodynamic Research - NASA Technical Paper 3675. Technical report, NASA, 1997.
- [9] C Eschmann, C Kuo, and C Boller. Unmanned Aircraft Systems for Remote Building Inspection and Monitoring. *Proceedings of the 6th European Workshop on Structural Health Monitoring, July 3-6, 2012, Dresden, Germany*, pages 1–8, 2012.
- [10] Federal Aviation Administration. Helicopter Components, Sections, and Systems. In *Helicopter Instructor's Handbook*, chapter Chapter 5, page 183. US Department of Transportation, Oklahoma, 2012.
- [11] Kenneth Warren Flanigan. Gas Powered Tip-Jet-Driven Tilt-Rotor Compound VTOL Aircraft, 2006.
- [12] Lee N Hager. Drive System for a Variable Diameter Tilt Rotor, 2000.
- [13] W. Johnson. Camrad a Comprehensive Analytical Model of Rotorcraft Aerodynamics and Dynamics. Technical report, NASA, California, 1980.
- [14] Adam Klaptocz, Adrien Briod, Jean-christophe Zufferey, and Dario Floreano. An Indoor Flying Platform with Collision Robustness and Self-Recovery. pages 3349–3354, 2010.
- [15] J. Gordon Leishman. Principles of Helicopter Aerodynamics. Cambridge Aerospace Series. Cambridge University Press, 2nd edition edition, 2006.
- [16] Andrew H Logan and Richard E Moore. Helicopter Antitorque System Using Circulation Control, 1980.
- [17] Teppo Luukkonen. Modelling and Control of Quadcopter. Journal of the American Society for Mass Spectrometry, 22(7):1134–45, 2011.
- [18] M.W. Mueller and R. D'Andrea. Stability and control of a quadrocopter despite the complete loss of one, two, or three propellers. In *Robotics and Automation (ICRA), 2014 IEEE International Conference* on, pages 45–52, May 2014.
- [19] Yogianandh Naidoo, R Stopforth, and Glen Bright. Rotor Aerodynamic Analysis of a Quadrotor for Thrust Critical Applications. Technical Report November, University of KwaZulu Natal, Durban, 2011.
- [20] Paul Y Oh, Michael Joyce, and Justin Gallagher. Designing an Aerial Robot for Hover-and-Stare Surveillance. *IEEE Advanced Robotics*, (12):303–308, 2005.
- [21] Kevin Sablan. Theory of flight.
- [22] La Young, Ew Aiken, Jl Johnson, J Andrews, J Klem, and R Demblewski. New concepts and perspectives on micro-rotorcraft and small autonomous rotary-wing vehicles. Technical report, US Army Aviation, 2002.

# High-level Rapid Prototyping of Graphical Models

Hendrik de Villiers<sup>1</sup>, Lynette van Zijl<sup>1</sup> and Thomas Niesler<sup>2</sup>

Abstract—Graphical models provide a flexible framework for capturing probabilistic relationships between large numbers of variables. Because of this flexibility, researchers usually compare a number of possible model topologies and parameter settings for their particular application. However, typically a choice exists between libraries in lower-level languages with faster inference, and libraries in higher-level languages with slower inference. Here we present a rapid-prototyping library for graphical models which combines flexible representation in a high-level language with dynamic compilation to lowlevel code for performing actual inference. We demonstrate the effectiveness of the resulting library on a set of benchmark images, comparing with a reference implementation in the C++ language. We show that the library facilitates the creation of domain-specific languages for quickly defining and testing graphical models.

# I. INTRODUCTION

Probabilistic graphical models (PGMs) [1], [2] enable reasoning about uncertain information in problem domains with complex structure, and yet simplify the capturing of this structure by concentrating on local relationships between the variables that describe the problem.

While the probabilistic relationships encoded in a graphical model are subjective in nature, once encoded, reasoning about these relationships proceeds in a quantitative fashion. So, subjective assumptions are made upfront in a principled way, and this is followed by objective inference given these assumptions.

Researchers employing PGMs typically experiment with a variety of prototype model structures and probabilistic relationships for purposes of identifying promising candidates. Therefore, the rapid prototyping of graphical models is key to accelerating progress within research programs that utilise them.

However, there are a number of competing needs for a rapid prototyping platform. Typically, a choice between these needs is necessary. On the one hand, low-level implementations can be fast at execution, yet are less easy to employ and customize. Alternatively, a more high-level implementation may provide high levels of introspection abilities and easy customization, but suffer in terms of speed.

We propose, implement and test a library CL-PGM (Common Lisp Probabilistic Graphical Models) for message-based inference in PGMs which manages to combine the ease and introspection abilities of high-level implementations, with a dynamic compilation architecture which produces efficient machine code for performing actual inference. In the subsequent sections, we provide a brief overview of graphical models, and discuss existing inference libraries with a particular focus on the OpenGM library [3]. Focussing on inference using Loopy Belief Propagation (LBP), we then present the proposed library architecture, and analyse our implementation using benchmark experiments, comparing with results obtained using the LBP implementation made available by OpenGM. The results show that, despite being implemented in a flexible, high-level language, the proposed library is comparable in execution speed and memory usage to the reference C++ implementation. In combination with the ability to easily define new domain specific languages for quick model definition, these results suggest that CL-PGM has potential as a test-bed for rapid prototyping of graphical models.

#### II. GRAPHICAL MODELS: AN OVERVIEW

Graphical models are concerned with probabilistic models where the joint distribution over the variables of interest factorize as

$$p(\mathcal{X}) = \frac{1}{Z} \prod_{i=1}^{N} \phi_i(\mathcal{X}_i).$$
(1)

Here,  $\mathcal{X}$  is the set of all random variables of interest, and  $\phi_i(\mathcal{X}_i)$  are a set of non-negative potential functions over subsets  $\mathcal{X}_i$  of  $\mathcal{X}$ . The constant Z is a normalization constant such that the integral of  $p(\mathcal{X})$  over all the variables in  $\mathcal{X}$  is one, which is necessary for valid probability densities.

The potentials  $\phi_i(\mathcal{X}_i)$  specify local relationships between the subset variables  $\mathcal{X}_i$ , governing how favourable local settings of these variables are. If  $\phi_i(\mathcal{X}_i)$  is large for a particular setting of the variables  $\mathcal{X}_i$ , then the setting is locally favourable.

However, local favourability does not imply global favourability. In general, we must perform inference to find the total effect of all  $\phi_i(\mathcal{X}_i)$  taken together.

A popular, powerful class of algorithms for resolving these questions work by passing messages in graph structures. These graphs are based on the independence relationships implied by Equation 1. A number of graph representation possibilities exist, one particularly flexible representation is the Cluster Graph [2].

While the origin of Cluster Graphs is beyond the scope of this article, describing their key properties and intuitive meaning is sufficient for the remaining sections.

Cluster graphs allow one to represent connections between variables of the joint distribution. For example, the joint distribution

<sup>&</sup>lt;sup>1</sup>Hendrik de Villiers and Lynette van Zijl are with the Computer Science Division, Stellenbosch University, South Africa {hdevil,lynette}@cs.sun.ac.za

<sup>&</sup>lt;sup>2</sup>Thomas Niesler is with the Department of Electrical and lectronic Engineering, Stellenbosch University, trn@sun.ac.za



Fig. 1: Example Cluster Graph

$$p(A, B, C, D) = \phi_1(A, B)\phi_2(A, B, C)\phi_3(B, C, D)\phi_4(C)$$
(2)

can be represented using a cluster graph as shown in Figure 1. Elliptical nodes are referred to as clusters, which have associated with them a set of random variables  $\kappa_n$ , referred to as their cluster variables. Each potential in the joint must be associated with one, and only one, cluster. All the variables bound to the potential must be present in the cluster as well.

In message passing algorithms, clusters communicate with each other by means of the edges between them, forming a consensus regarding the probability assigned to different settings of the random variables associated with them. On the edges between clusters are rectangular nodes labelled with a subset of the intersection of the variables in the clusters connected by the edge. This label is called a sepset. Information communicated over an edge during message passing pertains only to those variables in the sepset. There are two messages on each edge in the cluster graph, one forward, one backward.

To illustrate message passing, we consider calculating the message  $\mu_{2\rightarrow3}$  between the clusters ABC and BCD for Shafer-Shenoy sum-product message passing [1]. In order to calculate a message, we combine the potentials at the source node ABC (here just  $\phi_2$ ) with the incoming messages from other nodes (except the target node BCD). The incoming messages in this case are  $\mu_{1\rightarrow2}$  and  $\mu_{4\rightarrow2}$ . Finally, only information relevant to the sepset BC may be communicated over the edge, so we must marginalize away A. The resulting message then has the form

$$\mu_{2\to 3}(B,C) = \sum_{A} \phi_4(A,B,C) \mu_{1\to 2}(A,B) \mu_{4\to 2}(C).$$
(3)

One can view clusters as repositories of beliefs surrounding their cluster variables. The potentials represent prior beliefs about the cluster variables. Messages from other clusters update the beliefs about the cluster variables.

The CL-PGM library implements message passing of this kind in general cluster graphs. Before discussing the library architecture, we first discuss existing reference implementations that form part of the OpenGM library.

# III. OPENGM

OpenGM [3], [4] is a C++ library offering an extensive variety of inference strategies, with the explicit goal of providing a benchmark framework for comparing different inference approaches. It provides wrappers for other popular inference libraries, such as libDAI [5] and MRF-lib [6].

CL-PGM grew out of the need for interactive experimentation with graphical model architectures in a highlevel language. While the OpenGM library supports the creation of extensions and provides interfaces for the Python and MATLAB languages, the question arises whether an implementation crafted entirely within a high-level language can perform comparably.

Another motivation for CL-PGM's development was that OpenGM does not support general cluster graphs. Instead, factor graphs, which are special cases of cluster graphs, are employed for model specification. The use of CL-PGM's facilities for general cluster graph handling are demonstrated in [7], where the library was successfully employed as part of an automatic colour calibration system. However, for the present article, we will focus on comparing OpenGM and CL-PGM's performance. Therefore, the analyses presented here are limited to benchmark models represented using factor graphs.

In the following section, we discuss the architectural choices that allow CL-PGM to perform efficient inference, despite its implementation in a high-level, weakly typed, garbage-collected language.

#### **IV. SYSTEM ARCHITECTURE**

In this section, the architecture of the CL-PGM library is described. We will discuss the choice of target language, the library structure and the manner in which message passing is implemented to enable inference to be performed efficiently.

#### A. Target language

The choice of implementation language is critical to successfully creating a library which is flexible and easy to use, yet allows inference to be performed at speeds and with memory usage comparable to low-level implementations.

Towards this end, Common Lisp [8], [9] was chosen, in particular the Steel Bank Common Lisp (SBCL) [10] compiler. This platform is attractive for a number of reasons:

1) High-level Language Features: Common Lisp provides rich, object-oriented semantics. Functions are first class objects, and it is simple to specify anonymous functions (lambdas). As with languages like Python, Common Lisp is weakly typed (variables and parameters do not need to be declared as being of a certain type). Run-time type checking and array bounds checking are performed by default. Furthermore, the language features automated garbage collection, facilitating the reliable disposal of complex, interconnected data structures such as those that might be used during the construction of PGMs. Finally, the style of development in Lisp is interactive, programmers typically interact with a live Lisp image, similar to languages such as MATLAB.
```
(defun unoptimized-sum (n)
  (let ((sum 0))
    (dotimes (x n) (incf sum x))
    sum))
 disassembly for UNOPTIMIZED-SUM
;
 Size: 143 bytes
  ... skipping 15 lines
;
  GENERIC-+
       1C8:
                 MOV R11D, 536871408
       1CE:
                 CALL R11
  ... skipping 3 lines
;
  GENERIC-+
;
                 MOV R11D, 536871408
       1DE:
       1E4:
                 CALL R11
     skipping 3 lines
;
  . . .
  GENERIC-<
;
       1F3:
                 MOV ECX, 536871925
;
       1F8:
                 CALL RCX
     skipping 5 lines
       205:
                 RET
;
  ... skipping 11 lines of error handling
;
(defun optimized-sum (n)
  (declare (optimize (speed 3)
                       (safety 0)
                       (debug 0)))
  (let ((sum 0))
    (declare (type fixnum n sum))
    (dotimes (x n) (incf sum x))
    sum))
 disassembly for OPTIMIZED-SUM
;
;
  Size: 32 bytes
  1F6D3842:
                 XOR EDX, EDX
;
        44:
                 XOR ECX, ECX
;
         46:
                 JMP L1
;
                 NOP
        48:
;
        50: LO: ADD RDX, RCX
;
        53:
                 ADD RCX, 2
        57: L1: CMP RCX, RBX
;
        5A:
                 JL LO
;
        5C:
                 MOV RSP, RBP
;
;
        5F:
                 CLC
        60:
                 POP RBP
        61:
                 RET
```

Fig. 2: Disassembled machine code of a simple summation loop for optimized and unoptimized versions generated by SBCL for the x86-64 architecture. Without type declarations, generic, slow, versions of + and < are called. When optimized, the ADD and CMP machine instructions are used directly. It is similarly possible to declare array types, leading to fast access without bounds checking.

These features facilitate rapid development, relieving researchers of the unnecessary effort in writing routine "bookkeeping" code, assisting debugging, and allowing customization and interaction with the PGM representation at run-time.

2) Low-level Language Performance: While high-level operations facilitate the *definition* of new PGMs, inference still has to be performed efficiently. Common Lisp implementations such as SBCL support a variety of mechanisms allowing one to write programs which compile to efficient machine code.

Common Lisp features optional type declarations, which can be used to perform automatic type checking, but also to provide hints for compilers to generate fast low-level code.



Fig. 3: Library Architecture. Top layer: High-level representation as a cluster graph. This representation may optionally be discarded once the messages and message schedule have been created and compiled. Middle layer: Message schedule consisting of rapid triggering and committing of messages in a desired order. Lower layer: Message data and associated optimized code for calculating new versions when triggered. Dotted rectangles indicate that messages are double-buffered (each message must be "committed" after calculation).

Figure 2 gives a basic illustration of how types and levels of optimization are declared, and the dramatic effect this can have on the machine code produced. This flexibility allows Common Lisp to, by default, act as a high-level language, but allows speed critical sections of code to be optimized according to compiler hints.

The compiler itself forms part of the Lisp runtime, and so code may be defined and compiled on demand at runtime. This includes any anonymous functions, which forms the basis of CL-PGM's message passing architecture.

3) Metaprogramming Facilities: Lisp features strong facilities for metaprogramming. That is, programs in Lisp are actually data structures that can be manipulated or written by other Lisp code [9]. This allows the process of writing low-level code to be automated, so that library users need only interact with the high-level aspects of the language.

The macro system made available by Common Lisp is powerful and flexible, as the full language runtime is available while code is being automatically generated, in contrast with macro/template systems of languages such as C++, which are more limited in the kinds of code transformations they can perform.

#### B. Library Overview

Figure 3 illustrates the basic structure of the representation employed by the library. We combine a high-level description layer (the top layer) in terms of cluster graphs, with a lowlevel message passing layer (the bottom layer) controlled by a middle scheduling layer.

The cluster graph representation layer supports a suite of operations for defining and manipulating cluster graphs, random variables and potentials. This level provides introspection functions allowing the structure of the network to be easily available for client applications. The message passing and scheduling layers are generated from the representation layer. We adopt a message-centric architecture for the low-level inference portion, where each message is accompanied by an optimized calculator function that may be triggered by the scheduling layer. The calculator functions avoid the overhead of object-oriented approaches by being simple, parameterless, anonymous functions which maintain pointers to the message potential tables. Messages are also double buffered, each having two tables, where one stores a newly calculated version of the message until it is committed.

By design, the high-level description of the network may be discarded to save memory, if needed. This leaves the lowlevel description, which is faster and more compact.

#### C. Message Calculators

The inference engine is distributed across a set of message calculators (bottom layer in Figure 3), two for each edge in the cluster graph. As shown, each of these message calculators maintains two arrays containing message data.

Each of the message calculators is custom compiled based on the potentials and messages that they take as input, as well as the variables bound to these potentials and messages. The code is generated and compiled dynamically at runtime before the beginning of the message passing schedule. Because the code is generated algorithmically, type and optimization declarations are added automatically in order that efficient machine code is generated.

Consider the following minimization

$$\min_{C,D} \phi_1(A, B = 3, C) \phi_2(C, D, A) \phi_3(C, B = 3, D).$$

A number of optimization may be performed to implement this message calculation. Firstly, one can eliminate inferencetime interaction with the potential objects  $\phi_n$  by obtaining references to their underlying table arrays. Using the variable objects  $A, \ldots, D$ , one can determine which loops need to be established and over what domain the loops need to be performed (in this case over C and D). Constants such as B may explicitly substituted in the generated code. Lookups within the potential tables can be reduced to simple array access statements, as the ordering of the variables is available at message setup time. The types of each quantity may be declared, and bounds checking of arrays turned off.

Combined, these optimizations allow rapid computation of new versions of messages at inference-time. However, for very large networks with large groups of similar messages, compiling a pair of custom message calculators for each edge in the cluster graph is wasteful in terms of memory and setup time (calls to the compiler itself are relatively slow). To address this, we introduce a caching system which identifies whether a given message may reuse the calculation code of an existing message calculator. If this is the case, a version of the existing calculator is used, bound to the data of the newly set up message. This enables us to avoid excessive setup times in graphical models such as Markov Random Fields which exhibit a high number of messages similar in structure.

The CL-PGM library currently supports damped loopy belief propagation, where the summation  $\Sigma$  and product  $\Pi$  operators may also be any operators which together form a semiring. In particular, this means that sum-product messages ( $\Sigma$ ,  $\Pi$ ), max-product messages (max,  $\Pi$ ) and minsum messages (min,  $\Sigma$ ) are supported.

Calculated messages may be committed as they are or, to help suppress oscillations in message beliefs, they may be blended by some constant factor with the old version of the message (see "damping" Koller [2]).

#### V. EXPERIMENTAL SETUP

We test the efficiency of message passing in CL-PGM by comparing its Loopy Belief Propagation implementation with that of OpenGM. A set of benchmarks accompanying OpenGM using its HDF5-based format [3] will be employed for this purpose. Routines allowing the necessary benchmarks to be loaded into CL-PGM were written. The benchmark factor graph representations were converted to the corresponding Bethe cluster graph representations [2] to match CL-PGMs internal representation.

Two image processing tasks were considered utilizing variants of Markov Random Fields. These benchmarks are available on the OpenGM site [11]: In-painting and Color Segmentation.

The first tasks were two in-painting tasks, one shown in Figure 6 (the left-most side is the initial image). Two kinds of connectivity between neighbouring pixels were considered, 4-connected and 8-connected neighbourhoods. The other image processing tasks were 21 image segmentation tasks, with again a mixture of 4-connected and 8-connected neighbourhoods. Inference was performed using damping set at 0.5, with messages being updated and then committed in parallel.

#### VI. RESULTS AND CONCLUSION

Figures 6, 7, 8 illustrate the results of the inference process (as performed by CL-PGM). These results were similar to those obtained with OpenGM.

We consider three key performance metrics: the speed of inference, the amount of memory consumed during inference, and the goodness of the final solutions (how low their energy is). If the joint distribution is  $p(\mathcal{X}) = \prod_i \psi(\mathcal{X}_i)$ , then the energy of a solution  $\mathcal{X}'$  is  $\sum_i \log \psi(\mathcal{X}'_i)$ . To measure these in a way that is independent of the particular machine on which execution takes place, we report the ratios of these quantities for each of the test cases, combined as single histograms. Note we have focussed on inference time behaviour. However, the initial set up time and memory requirements of CL-PGM tended to be higher because of the compilation overhead.

The top subfigure in Figure 5 is a histogram of the ratios between the minimum (best) energies of the solutions for each task obtained by CL-PGM and OpenGM respectively.

```
(make-factor-graph
; Declare variables and domains
             '(nil t)) ; nil=false, t=true
((burglar
             '(nil t) t) ; observed as t
 (alarm
  (earthquake '(nil t))
             '(nil t) t)); observed as t
 (radio
; Declare potentials
((burglar-prior (burglar)
                 #(0.99 0.01))
                  1% chance
  (earthquake-prior (earthquake)
                    #(0.9999 0.0001))
                    : 0.01% chance
 (radio-pot (radio earthquake)
   (tabulate
      (probability-of radio
                 (if earthquake 1.0
                                (0.0))))
  (alarm-pot (alarm burglar earthquake)
    (tabulate
      (probability-of alarm
        (cond ((and earthquake burglar) 0.999)
               (burglar 0.99)
               (earthquake 0.99)
               (t 0.001))))))))
```

Fig. 4: Example of a small domain specific language for defining factor graphs. Comments begin with semicolons.

It can be seen that CL-PGM and OpenGM produce solutions of similar quality (the ratios are all close to 1).

The middle subfigure in Figure 5 illustrates the ratio between the average step time taken during inference for CL-PGM and OpenGM respectively. It can be seen that for approximately 61% of the tests, the ratio was close to unity, meaning that OpenGM and CL-PGM were performing inference steps at around the same speed (with CL-PGM slightly outperforming OpenGM in 50% of cases). In the other 39% of cases, OpenGM was performing inference steps roughly twice as fast as CL-PGM. The mean of the distribution of ratios is 1.38, showing that CL-PGM operates at speeds comparable to that of OpenGM, suggesting that the dynamic compilation strategy has been successful.

The lower subfigure in Figure 5 shows the ratio of memory consumption during inference for both systems. For each test case, the ratios are close to one, slightly favouring OpenGM in some cases with up to 20% larger memory consumption in CL-PGM (ratio 1.2). This shows that inference time memory consumption is roughly the same for both libraries. Note however that it was found that CL-PGM uses approximately the same amount of memory in addition during compilation.

We have further tested this by successfully employing CL-PGM for automatic colour model calibration for a glovebased hand tracker [12], [7].

However, while raw inference performance is important, one may now further abstract the library functionality by defining domain specific languages for describing graphical models. We have done so for both Factor Graphs and Cluster Graphs. A small example is shown in Figure 4, where the well-known Burglar network [1] is implemented using a mini-language that is transformed by the Common Lisp macro system into a Bethe cluster graph for use in inference.

These results demonstrate the viability of implementing



Fig. 5: The ratio of CL-PGM to OpenGM's average time taken per inference step, ratio of the lowest energies obtained, and the average amount of inference-time memory used. Test cases with ratios less than 1 favour CL-PGM, ratios more than 1 favour OpenGM.

inference frameworks in high-level languages, without excessive loss of performance, if sufficient meta-programming facilities are available to enable automatic compilation to low-level code.

#### ACKNOWLEDGMENT

The authors wish to thank the Wilhelm Frank Trust, Telkom SA and the National Research Foundation of South Africa (grant UID 71926) for providing financial support during the course of this research.

#### REFERENCES

- [1] D. Barber, *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.
- [2] D. Koller and N. Friedman, Probabilistic graphical models: Principles and techniques. MIT Press, 2009.
- [3] J. H. Kappes, B. Andres, F. a. Hamprecht, C. Shnorr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, T. Kroger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother, "A Comparative Study of Modern Inference Techniques for Structured Discrete Energy Minimization Problems," *International Journal of Computer Vision*, 2015. [Online]. Available: http://arxiv.org/abs/1404.0533



Fig. 6: In-painting example with 8-connected pixel neighbourhood constraints. Pixels outside the central circle in the first frame are observed, and stay constant during all steps. The remaining central pixels are updated through message passing, starting with a uniform initialization. From left to right: Inference steps 1, 30, 50, 100 and 300 respectively.



Fig. 7: Color segmentation results. The left two columns are false colour and segmented versions of the image after the first inference step. The right two columns are taken after inference has completed.



Fig. 8: Color segmentation results. The original image (left) is segmented and progressively refined towards the rightmost image.

- [4] J. H. Kappes, B. Andres, F. a. Hamprecht, C. Schnorr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, J. Lellmann, N. Komodakis, and C. Rother, "A Comparative Study of Modern Inference Techniques for Discrete Energy Minimization Problems," 2013 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1328–1335, Jun. 2013. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6619019
- [5] J. M. Mooij, "libDAI : A Free and Open Source C ++ Library for Discrete Approximate Inference in Graphical Models," *Journal of Machine Learning Research*, vol. 11, pp. 2169–2173, 2010.
- [6] R. Szeliski, R. Zabih, S. Member, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother, "A Comparative Study of Energy Minimization Methods for Markov Random Fields with Smoothness-Based Priors," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 30, no. 6, pp. 1068–1080, 2008.
- [7] H. A. C. de Villiers, "A vision-based South African Sign Language tutor," Dissertation, Stellenbosch University, 2014. [Online]. Available: http://hdl.handle.net/10019.1/86333
- [8] G. L. Steele, *Common Lisp: The Language*, 2nd ed. Digital Press, 1990.
- [9] P. Seibel, Practical Common Lisp. Apress, 2005.
- [10] "Steel Bank Common Lisp." [Online]. Available: http://www.sbcl.org/
- [11] "OpenGM benchmark page." [Online]. Available: http://hci.iwr.uniheidelberg.de/opengm2/?10=benchmark
- [12] H. de Villiers, L. van Zijl, and T. Niesler, "Visionbased hand pose estimation through similarity search using the Earth Mover's Distance," *IET Computer Vision*, vol. 6, no. 4, pp. 285–295, 2012. [Online]. Available: http://digitallibrary.theiet.org/content/journals/10.1049/iet-cvi.2011.0128

# Unsupervised acoustic model training: comparing South African English and isiZulu

*Neil Kleynhans*<sup>1</sup>, *Febe de Wet*<sup>2</sup> and *Etienne Barnard*<sup>3</sup>

<sup>1,3</sup> Multilingual Speech Technologies, North-West University, Vanderbijlpark, South Africa
 <sup>2</sup> Human Language Technologies Research Group, Meraka Institute, CSIR, South Africa
 <sup>2</sup>Department of Electrical and Electronic Engineering, Stellenbosch University, South Africa

{ntkleynhans,fdwet,etienne.barnard}@gmail.com

#### Abstract

Large amounts of untranscribed audio data are generated every day. These audio resources can be used to develop robust acoustic models that can be used in a variety of speech-based systems. Manually transcribing this data is resource intensive and requires funding, time and expertise. Lightly-supervised training techniques, however, provide a means to rapidly transcribe audio, thus reducing the initial resource investment to begin the modelling process.

Our findings suggest that the lightly-supervised training technique works well for English but when moving to an agglutinative language, such as isiZulu, the process fails to achieve the performance seen for English. Additionally, phone-based performances are significantly worse when compared to an approach using word-based language models. These results indicate a strong dependence on large or well-matched text resources for lightly-supervised training techniques.

Index Terms: lightly-supervised training, unsupervised training, automatic transcription generation, audio harvesting, English, isiZulu

#### 1. Introduction

Vast amounts of audio data are created on a daily basis. Typical sources are radio / television broadcasts, podcasts and lectures. Very few of these audio corpora have corresponding orthographic or other transcriptions. A particularly interesting scenario where large amounts of audio data are created and where text transcriptions would be of great benefit are call-centre environments. Access to text representations of the audio would aid in swiftly analysing the data and making necessary adjustments where appropriate.

Manually transcribing audio data is a resource intensive process requiring disproportionate amounts of money and time. The time to produce a transcription depends on transcriber expertise and required accuracy of transcription: Approximate transcriptions can be generated at 3 to 5 times real time while for highly accurate transcriptions the time considerably increases to 50 times real time [1]. This turnaround time is often too long to make business sense, considering the amount of audio data collected in a day, and so automatic means become the only feasible route.

An automatic solution would require an automatic speech recognition (ASR) system at its core, but in general one would not have access to matching in-domain audio and text data. To sidestep this dilemma, the "lightly supervised" acoustic model (AM) training approach [2] provides a mechanism to develop and constantly refine AMs needed by an ASR system. This approach, however, requires approximate transcriptions (for instance, closed-caption transcriptions are frequently employed). Another approach is "unsupervised" AM training [2, 3] which follows the same broad steps as the lightly-supervised approach but does not make use of approximate transcriptions. What makes these approaches attractive is the minimal initial resource investment. An added incentive is that the AMs are trained on in-domain data, which removes the mismatch between the audio data and the AMs.

In this paper we investigate factors related to the resources required to start and maintain the automatic harvesting of untranscribed audio data using an unsupervised AM training approach. Specifically, we investigate the scenario for the resources-constrained South African English (SAE) and isiZulu languages.

#### 2. Background

Lightly-supervised and unsupervised acoustic model training have been applied in many different scenarios [2, 3, 4, 5, 6, 7, 8, 9]. The basic steps of the iterative algorithm are [2] the following:

- Partition the audio data into homogeneous portions based on characteristics such as channel, environment or speaker.
- Normalise all approximate transcriptions (if available) and produce appropriate phonetic pronunciations.
- Develop seed acoustic models by manually transcribing a small portion of the in-domain audio data or using existing AMs.
- Automatically transcribe the raw audio data using welltrained or biased language models (LM) or word-graphs trained on the approximate transcriptions (e.g. closed captions).
- As an optional step, use acoustic models to align the audio and the available approximate transcriptions. Remove audio data where the alignment and automatic transcription disagree excessively.
- Use the newly transcribed data to re-train AMs and repeat the process.

If the penultimate step is not implemented then the lightlysupervised approach collapses to an unsupervised training approach.

Lamel *et al.* [2] investigated the minimal requirements needed to bootstrap the lightly-supervised and unsupervised

processes. AMs trained on 10 minutes of audio data and a LM developed on 1.8 M word corpus (40k lexicon) were sufficient to initiate the process. On 200 hours of raw audio data, using the Topic Detection and Tracking (TDT-2) English corpus, the minimal system obtained a final word-error-rate (WER) of 28.8%, which was significantly worse compared to 18% achieved by a system utilising a LM trained on text corpora (Hub4 and TDT corpora) - but these corpora contain orders of magnitude more text. In a completely unsupervised approach, the minimal system achieved a WER of 37.4% on 140 hours of raw audio. Similarly, the same unsupervised approach was applied to train AMs on Portuguese broadcast news audio [4]. The 3.5 hour trained AM produced a WER of 42.6% while AMs trained on automatically transcribed 30 hours delivered a WER of 39.1%. The LMs were trained on news-related text corpora containing 72.6M words.

Novotney *et. al.* [5] experimented with limited amounts of labelled audio and text data in their unsupervised AM training investigations. They focused on the Fisher corpus data, containing telephone-quality conversational speech and had access to a text corpus containing 1.1 billion words. Interestingly, their findings suggest weaker LMs do not severely impact the unsupervised training of AMs and have a greater impact when decoding the actual evaluation set.

To improve their transcribing system Nguyen and Xiang [6] added 702 hours of audio selected from a 1 400 hour audio data set using the lightly-supervised procedure. The audio data was selected from the TDT English corpora – TDT2, TDT3 and TDT4. The initial acoustic models were trained on 141 hours of audio data and the subset-specific (depending on data set) LMs were trained by interpolating from a 360M word common LM. The baseline system WER of 12.1% was reduced to 10.1% using the 702 hours of added training data.

Gales *et. al.* [8] made use of lightly-supervised training to automatically transcribe audio data from the TDT and 2003 BN collection corpora. The biased LMs were trained on broadcast news text as well as closed captions. The LM weighting was similar to that of Hazen [1] - 90% closed-caption text and 10% general broadcast news.

Chan and Woodland [7] applied lightly-supervised training to 500 hours TDT2 and 300 hours TDT4 corpora. The LMs were trained on text sourced from the closed captions as well as closely related text. The entire text corpus consisted of approximately a billion words. The out-of-vocabulary rates were 0.68% and 0.47% for the different corpora. Again, a biased LM was used with interpolation weights of 0.92 and 0.90, respectively.

Gollan *et. al.* [3] utilised unsupervised training to improve upon baseline AMs trained on 100 hours of manually transcribed audio data selected from English-only European Parliament Plenary Session speeches. Adding 180 hours of automatically transcribed audio improved the WERs from 10.4% to 9.6%.

Davel *et. al.* [9] showed that lightly-supervised automatic harvesting for ASR resource creation in a resource-scarce environment does not require well-trained LMs. In their approach, a phone-based ASR system was used to automatically generate transcriptions, for roughly 100 hours of SAE radio broadcast audio data, using a flat-phone task grammar. The seed models were initially developed on US English data and gradually replaced by the in-domain SAE dialect. Data filtering was achieved by using a garbage model that absorbed badly aligned audio portions.

Gelas *et. al.* [10] made use of a Swahili ASR system to aid in speeding up the task of manually transcribing a 12 hour

audio portion of a 200 hour Swahili web broadcast news speech corpus. The initial ASR system was trained on 3.5 hour read speech Swahili corpus. The procedure used an ASR to automatically transcribe a 2 hour portion of audio. These transcriptions were manually corrected. After the correction process, the newly transcribed audio data was used to increase the amount of training data used to train the new AMs. The process was repeated until 12 hours were transcribed. Utilising the ASR system to automatically transcribe the data reduced the manual correction time from an initial 40 hours to a final 15 hours. The LM was trained on a text corpus which contained 28 M words and had a 65k lexicon.

Previous investigations suggest that the unsupervised AM training approach does not require vast resources to begin the harvesting process. As few as 10 minutes of labelled audio data to train AMs and 100k - 1M words to train language models. Some approaches do not require LMs – Davel *et. al.* [9] – but approximate transcriptions were available for data filtering. If LMs are used, however, the text corpora are quite well matched to the domain which in a resource-constrained environment will not be easy to access or develop.

In this study we therefore investigate:

- unsupervised AM training without the aid of a language model phone decodes only ,
- the usefulness of language models trained on unrelated text corpora, and,
- the effect of text corpus size used to train N-gram LMs.

#### 3. Method

#### 3.1. Corpora

The NCHLT corpus is a read-speech corpus containing highbandwidth audio data and transcriptions thereof for all eleven South African languages [11]. Mobile devices were used to collect the audio data. The transcriptions contain short sentences and were derived from large text corpora in order to attain coverage of the most common triphones of the target language. For our unsupervised AM training investigations we limited ourselves to using the English and isiZulu sub-corpora.

#### 3.1.1. NCHLT English

The English NCHLT sub-corpus contains audio data collected from 210 different speakers. There are a total of 77 412 utterances with each speaker contributing roughly 500 utterances. Table 1 shows the duration in hours, the number of speakers and utterance amount for the training and evaluation data sets for the NCHLT English sub-corpus

Table 1: The duration, amount of speakers and number of utterances for the NCHLT English sub-corpus.

Data Set	Duration (Hours)	# speakers	# utterances
Training	54.19	202	74180
Evaluation	2.42	8	3232

There is a total of 223 561 tokens and a lexicon of 8 350 words for the entire corpus. The training set contains 214 192 tokens in total and a lexicon of 8 328 words, while the evaluation set contains a total of 9 369 tokens and a lexicon of 3 627 words. The out-of-vocabulary (OOV) rate between the training and evaluation set is 0.61%.

#### 3.1.2. NCHLT isiZulu

Similar to the English sub-corpus, the isiZulu NCHLT subcorpus contains audio data collected from 210 different speakers. There are 44 673 utterances in total with each speaker contributing roughly 500 utterances. Table 2 shows the duration in hours, the numer of speakers and utterance amount for the training and evaluation data sets for the NCHLT isiZulu sub-corpus

Table 2: *The duration, amount of speakers and number of utterances for the NCHLT isiZulu sub-corpus.* 

Data Set	Duration (Hours)	# speakers	# utterances
Training	52.23	202	41871
Evaluation	4.02	8	2802

There is a total of 133 480 tokens and a lexicon of 25 651 words for the entire corpus. The training set contains 125 028 tokens in total and a lexicon of 25 231 words, while the evaluation set contains a total of 8 452 tokens and a lexicon of 5 189 words. The out-of-vocabulary rate between the training and evaluation set is 8.1%.

#### 3.2. Pronunciation Modelling

The pronunciation dictionaries for the NCHLT sub-corpora were sourced from previous work as outlined in Davel and Martirosian [12].

The English pronunciation dictionary contained 15 000 unique entries and a phone set of 43 phones. Phonetisaurus [13] was used to perform grapheme-to-phoneme (G2P) prediction for words not found in the seed pronunciation dictionary. Phonetisaurus implements a WFST-driven G2P framework that can rapidly develop high quality G2P or P2G systems. The English NCHLT text required 3 966 G2P predictions.

For isiZulu a more elaborate approach was followed. For isiZulu words only G2P prediction was performed using the default&refine algorithm proposed in [14], while for codeswitched English words, the above Phonetisaurus G2P prediction was used. To identify English words a simple N-gram textbased language identification was implemented. The MIT language modelling toolkit [15] was used to build 3-gram back-off LMs for both English and isiZulu. The training word sets were extracted from the seed pronunciation dictionaries - the English word set had 15 000 words in total and isiZulu had a total of 15 404 words. A word was classified based on the perplexity score. Once all words with missing pronunciations were predicted the English phone set was mapped to the isiZulu phone set using manual rules. Lastly, the isiZulu phone set was further mapped using the MultiPron rules [16] which resulted in 32 phones in total.

#### 3.3. ASR system

The speech recognition system development follows a similar structure to that described in Kim *et. al.* [17]. The audio data was converted to Perceptual Linear Prediction (PLP) coefficients. The 52 dimensional feature vector was created by appending the first, second and third derivatives to the 13 static coefficients (including the 0'th component). Corpus-wide mean and variance normalisation was applied.

AMs were developed by following an iterative training scheme. Firstly, 32-mixture context-independent (CI) AMs were trained and used to produce state aligns for the CI AMs trained in the initial development of cross-word triphone context-dependent (CD) AMs. Once the CD AMs were trained the process was repeated and the previous AMs were used to produce all state alignments before the model mixture incrementing phase. The process was repeated twice for all experiments.

All Hidden Markov Models (HMM) employed a three state left-to-right structure. Each CD HMM's state contained eight mixture diagonal covariance Gaussian models. A questionbased tying scheme was followed to create a tied-state data sharing system [18] - where any context-dependent triphone having the same central context could be tied together.

Once the CD AM development was completed, Heteroscedastic Linear Discriminant Analysis (HLDA) was applied to reduce the 52-dimensional PLP feature vectors to a dimension of 39. A global transform was used for the estimation – a single class for all the triphones. After estimating the HLDA transform, the CD AMs' parameters were updated. It was found that allowing the variances to be updated resulted in a large percentage of floored variances. Therefore, only the weights and mean parameters were updated. Two update iterations were performed.

Lastly, Speaker Adaptive Training (SAT) was applied using Constrained Maximum Likelihood Linear Regression (CM-LLR) transformations. The same HLDA global transform was used and the CD AMs were updated twice – only weights and means.

The decoding task was a two-step process. The HLDA CD AMs were used to automatically generate transcriptions and a speaker-based CMLLR transform estimated. Then the CMLLR was applied on the second decoding pass.

#### 3.4. Language Models

To investigate the effect of developing LMs on mismatched text corpora, two alternate sources of text unrelated to the NCHLT corpora were used. Before training the LM, the text had to be normalised. This involved,

- Removing punctuation marks.
- Converting numbers to written form.
- Converting characters to lower-case.

Once normalised, the MIT-LM toolkit was used to develop the LMs. Only back-off bigram LMs were created due to limitations of HVite (the HTK decoder). For probability smoothing, fixed Kneser-Ney smoothing was applied.

#### 3.4.1. English

The English LM was developed on a 1.6M word text corpus. The text forms part of the 109M word South African Broadcast News (SABN) text corpus [19]. This corpus contains text extracted from a number of major South African newspapers. The text contain 1 692 929 tokens and a lexicon of 45 664 types. The OOV rate between the SABN and NCHLT text is 29.17%.

#### 3.4.2. isiZulu

The isiZulu LM was developed on a text corpus provided by the Centre for Text Technology (CText) [20]. The original corpus had 223 709 tokens but after applying text normalisation this number increased slightly to 234 216 (due to number expansion). The lexicon associated with the processed text was 38 869 types. The OOV between the CTEXT and NCHLT corpora was 71.59% – the OOV rate is far from ideal and may negatively influence the isiZulu results. Similar Zulu OOV rates have been seen in the investigation performed by Gales *et. al.* [21] and is common to morphological rich languages.

In addition to training a word-based LM, a syllable LM was also trained. Only words classified as isiZulu, using the text-based N-gram classification approach, were split into syllables. After splitting the words of the LM development text, there were 677 971 tokens and 7 759 types.

#### 3.5. Unsupervised training

To investigate the effectiveness of unsupervised AM training on SAE and isiZulu, the training sets of the NCHLT corpora were partitioned into a number of non-overlapping portions. The seed AMs were trained on a single hour selected at random from the entire training set and from 50 speakers. The transcriptions were used during the seed AM training which simulates the need for manually transcribing a portion of the audio if no other AMs or labelled data is available.

The remainder of the training data set was partitioned into smaller 3, 6, 12 and  $24^1$  hour sets of untranscribed audio. At each stage, the previously transcribed data sets, including the seed data, were pooled and used to develop new AMs. The current stage's untranscribed data was transcribed using the new AMs set.

To measure the progress of the unsupervised model training and accuracy of the models, phone-error-rates (PER) are reported on the evaluation set as well as the data set that was transcribed. PER are reported since, in the HTK model training recipes, only phone level representations are needed to train acoustic models.

Lastly, two methods of unsupervised acoustic model training were investigated. These are phone-based (flat phone grammars) and word-level LM-based approaches. Additionally, for isiZulu the syllable LMs are also investigated.

#### 4. Results

The PERs of the AMs developed through unsupervised training are reported. The accuracy is measured in terms of automatically transcribing the successive data portion (if for instance the AMs are trained on the seed plus three hours of data, the next six hour data set is viewed as a "testing" set) and the evaluation data set – the successive data portion set is labelled "Raw" and the evaluation set is labelled "Eval".

#### 4.1. English

Table 3 shows the AM PERs for increasing amounts of automatically transcribed data and using a flat phone-based grammar to harvest more data. Interestingly, when adding three hours of automatically transcribed data the error rates on the evaluation set decreases by more than 2% absolute; a 3% absolute decrease is seen on the raw six hour data set. After the +3 hour mark, however, PERs increase as more automatically transcribed data is used to train AMs.

Table 4 shows the PERs of AMs trained on automatically transcribed data and using a word LM model when decoding the data. The general trend is a decreased in PERs as more data is used to train the AMs, which is consistent with trends seen in literature. (The "Raw" +24 hr experiment was not reported in table 4, as "Raw" and "Eval" results are highly correlated and the same performance can be expected for the +24 hr "Raw" case).

Table 3: *The accuracy of the English acoustic models developed using a flat phone grammar approach.* 

Data Set	Raw	Eval
Seed (1 hr)	42.47	40.36
+ 3 hr	39.33	38.32
+ 6 hr	39.9	38.85
+ 12 hr	41.55	40.21

Table 4: The PERs of different English AMs trained on increasing portions of automatically transcribed data using a LM.

Data Set	Raw	Eval
Seed (1 hr)	25.26	23.52
+ 3 hr	21.33	20.39
+ 6 hr	19.11	18.74
+ 12 hr	14.66	14.73
+ 24 hr	-	13.98

#### 4.2. Text data dependency for English

Novotney *et. al.* [5] suggested that the size of the LM has a limited effect on the unsupervised training of AMs. To investigate this, we limited the amount of text used to train the English LMs. The text was limited to half and then a quarter of the full text. Table 5 shows the number of tokens, types and OOV rate for various sized text corpora used to train different LMs.

Table 5: Tokens, types and OOV of text used to develop LMs on full, half and quarter amounts of the English text corpus.

Percentage	Tokens	Types	OOV
100 %	1.69 M	45k	29%
50 %	846k	35k	34%
25 %	423k	26k	40%

Table 6 shows the AMs correctness and accuracies developed on increasing portions of automatically transcribed data using a LM trained of half the text corpus. As with the full text trained LM, all values increase as the amount of automatically transcribed data is used to train the AMs.

Table 7 shows the PERs obtained by using various AMs trained on automatically transcribed data and using a LM trained on a quarter of the full text corpus. Again, decreasing trends can be seen.

Considering the final evaluation results for AMs trained on all the acoustic data and the quarter, half and full sized LMs (17.36%, 16.78% and 13.98%), we can see a slight increase in accuracy as more text data is used to develop the LM. The drop in performance may also be attributed to the increase in OOV rates, observed for the LMs trained on less text data.

#### 4.3. isiZulu

Table 8 shows the performance of AMs trained on increasing amounts of automatically transcribed data using a flat phone decoding grammar. Besides a slight decrease in error rate for the raw set at the added three hour mark, the remaining PERs for both data sets steadily increase as more automatically transcribed data is added to the training pool.

Table 9 shows the PERs for various AMs trained on increas-

<sup>&</sup>lt;sup>1</sup>All the remaining data was around 24 hours in duration.

Table 6: The performance of English AMs used to automatically transcribe data using a LM developed on half the available text data.

Data Set	Raw	Eval
Seed (1 hr)	25.7	23.92
+ 3 hr	21.95	20.81
+ 6 hr	19.51	18.9
+ 12 hr	17.95	17.6
+ 24 hr	-	16.78

Table 7: *The performance of English AMs used to automatically transcribe data using a LM developed on a quarter of the available text data.* 

Data Set	Raw	Eval
Seed (1 hr)	25.7	24.99
+ 3 hr	22.4	21.49
+ 6 hr	20.07	19.44
+ 12 hr	18.69	18.31
+ 24 hr	-	17.36

ing amounts of harvested data and using a word LM during the decoding process. As with the flat phone approach, the same overall increasing trends are observed, however, the absolute performance values are somewhat lower.

Table 10 captures the performances of the AMs trained in an unsupervised manner while using a syllable LM during decoding. Again, there is a general increasing trend in the PERs as more automatically transcribed data is used to develop the AMs – except for the evaluation PER which decreases slightly for the added three hour mark. The PER values of the syllable approach are consistently better compared to the flat phone approach, but in general worse compared to the LM-based approach.

To try and rule out the possibility of poorly trained seed models, a different seed model trained on three hours of data was tried. Table 11 shows the PER percentages for AMs trained on increasing amounts of automatically transcribed data using a word LM during the decode cycle. Compared to the single hour seed model, the performance measures are slightly better but again the same increasing trend is observed as more data is added to the training pool.

#### 5. Conclusion

In this study we applied the well-known unsupervised acoustic model training scheme to resource-scarce South African English and isiZulu audio data. We investigated phone-based and word-based language models and, in addition, a syllable language model for isiZulu. The default seed acoustic model was trained on a single hour of manually transcribed data. The text corpora used to develop the language models were selected from unrelated sources which differed significantly in the OOV rates – 29% and 76% for English and isiZulu respectively. For English, we also experimented with the amount of text data used to train the language model.

From our results we may conclude:

- The unsupervised acoustic model training scheme performs well for SAE if a word-based LM is used.
- The phone-based approach, for English and isiZulu, did not achieve increasingly better results as more automatically transcribed data was added to the training pool.

Table 8: Unsupervised isiZulu AM training approach using a flat phone grammar.

Data Set	Raw	Eval
Seed (1 hr)	31.8	33.59
+ 3 hr	31.44	33.62
+ 6 hr	32.45	37.62
+ 12 hr	35.91	44.54

Table 9: Unsupervised isiZulu AM training approach using word LM.

Data Set	Raw	Eval
Seed (1 hr)	29.11	30.65
+ 3 hr	29.0	31.84
+ 6 hr	30.36	35.68

- For SAE, the word LM gave expected performances, according to the performance metrics, which confirms the importance of a language model when using unsupervised acoustic model training.
- The amount of text data used to develop the LM has a slight effect on the performances: even with relatively small amounts of text, successful unsupervised training was achieved in SAE, though increasing performance with more data added to the training pool was observed. The drop in absolute performance may be related to the increasing OOV rates.
- Based on all the isiZulu results, the application of the unsupervised acoustic model training approach was unsuccessful increasing amounts of automatically transcribed data produced poorer system accuracies. This is probably related to the high OOV rate of the isiZulu text corpus, which in turn results from the much larger vocabulary of a conjunctively written agglutinative language.
- Investigating the isiZulu results further showed: the correctness percentages for isiZulu increased, with increasing amounts of audio data but it was found that increasing the insertion penalty did not improve the accuracy values. This might suggest that the extremely high OOV rate severely limits the applicability of the unsupervised acoustic model training approach.

#### 6. Future Work

For future work, it would be informative to investigate whether unsupervised acoustic model training for isiZulu can be made to work with a sufficiently large text source, but also to understand whether the approaches that do not require such a text source can be adapted to succeed in this context.

Another unknown for the isiZulu investigation is the effect of out-of-language words. English does not suffer from this phenomenon and in the majority of cases is the "invader" language in isiZulu. Out-of-Language words are particularly bothersome with respect to pronunciation modelling, and our phonemapping approach is clearly a rough approximation in that case.

One possible approach to deal with the high isiZulu OOVs is to use a better syllabification approach. Our syllables did not yield any improvement over the word-based LM, but following the recent BABEL syllabification approach proposed by Davel

Table 10: Unsupervised isiZulu AM training approach using syllable LM.

Data Set	Raw	Eval
Seed (1 hr)	30.28	30.7
+ 3 hr	30.3	30.14
+ 6 hr	31.18	33.13
+ 12 hr	34.61	38.62

 Table 11: Unsupervised isiZulu AM training approach using word LM but starting with three hours of seed data.

Data Set	Raw	Eval
Seed (3 hr)	25.44	29.47
+ 3 hr	26.29	30.7
+ 6 hr	27.66	32.68

*et. al.* [22] may help to achieve successful automatic isiZulu harvesting.

#### 7. References

- T. J. Hazen, "Automatic alignment and error correction of human generated transcripts for long speech recordings," in *Proceedings of INTERSPEECH*. Pittsburgh, Pennsylvania, USA: ISCA, September 2006, pp. 1606–1609.
- [2] L. Lamel, J. L. Gauvain, and G. Adda, "Lightly supervised and unsupervised acoustic model training," *Computer Speech & Language*, vol. 16, no. 1, pp. 115–129, 2002.
- [3] C. Gollan, S. Hahn, R. Schlüter, and H. Ney, "An improved method for unsupervised training of LVCSR systems." in *Proceedings of INTERSPEECH*, Antwerp, Belgium, August 2007, pp. 2101–2104.
- [4] L. Lamel, J.-L. Gauvain, and G. Adda, "Unsupervised acoustic model training," in *Proceedings of the International Conference* on Acoustics, Speech and Signal Processing (ICASSP), vol. 1. Norwich, UK: IEEE, May 2002, pp. I-877–I-880.
- [5] S. Novotney, R. Schwartz, and J. Ma, "Unsupervised acoustic and language model training with small amounts of labelled data," in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP).* Taipei, Taiwan: IEEE, May 2009, pp. 4297–4300.
- [6] L. Nguyen and B. Xiang, "Light supervision in acoustic model training," in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 1. Montreal, Quebec, Canada: IEEE, May 2004, pp. I-185–I-188.
- [7] H. Chan and P. Woodland, "Improving broadcast news transcription by lightly supervised discriminative training," in *Proceedings* of the International Conference on Acoustics, Speech and Signal Processing (ICASSP), vol. 1. Montreal, Quebec, Canada: IEEE, May 2004, pp. I-737–I-740.
- [8] M. Gales, P. Woodland, H. Y. Chan, D. Mrva, R. Sinha, and S. E. Tranter, "Progress in the CU-HTK broadcast news transcription system," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 14, no. 5, pp. 1513–1525, 2006.
- [9] M. H. Davel, C. van Heerden, N. Kleynhans, and E. Barnard, "Efficient harvesting of Internet audio for resource-scarce ASR," in *Proceedings of INTERSPEECH*. Florence, Italy: ISCA, August 2011, pp. 3153–3156.
- [10] H. Gelas, L. Besacier, and F. Pellegrino, "Developments of Swahili resources for an automatic speech recognition system," in *SLTU-Workshop on Spoken Language Technologies for Under-Resourced Languages*, Cape Town, South Africa, May 2012.
- [11] N. J. de Vries, M. H. Davel, J. Badenhorst, W. D. Basson, F. de Wet, E. Barnard, and A. de Waal, "A smartphone-based ASR

data collection tool for under-resourced languages," *Speech communication*, vol. 56, pp. 119–131, 2014.

- [12] M. Davel and O. Martirosian, "Pronunciation dictionary development in resource-scarce environments," in *Proceedings of IN-TERSPEECH*, Brighton, United Kingdom, September 2009, pp. 2851–2854.
- [13] J. Novak, D. Yang, N. Minematsu, and K. Hirose, "Initial and evaluations of an open source WFST-based phoneticizer," *The University of Tokyo, Tokyo Institute of Technology.*
- [14] M. Davel and E. Barnard, "Pronunciation prediction with Default&Refine," *Computer Speech & Language*, vol. 22, no. 4, pp. 374–393, 2008.
- [15] B.-J. Hsu and J. Glass, "Iterative language model estimation: efficient data structure & algorithms," in *Proceedings of INTER-SPEECH*, vol. 8, Brisbane, Australia, September 2008, pp. 1–4.
- [16] N. Kleynhans, R. Molapo, and F. De Wet, "Acoustic model optimisation for a call routing system," in *Proceedings of the Annual Symposium of the Pattern Recognition Association of South Africa.* Pretoria, South Africa: PRASA, November 2012, pp. 165–172.
- [17] D. Kim, G. Evermann, T. Hain, D. Mrva, S. Tranter, L. Wang, and P. Woodland, "Recent advances in broadcast news transcription," in *Automatic Speech Recognition and Understanding*, 2003. *ASRU'03. 2003 IEEE Workshop on.* St. Thomas, U.S. Virgin Island: IEEE, November 2003, pp. 105–110.
- [18] S. J. Young, J. J. Odell, and P. C. Woodland, "Tree-based state tying for high accuracy acoustic modelling," in *Proceedings of the workshop on Human Language Technology*. Association for Computational Linguistics, 1994, pp. 307–312.
- [19] H. Kamper, F. de Wet, T. Hain, and T. Niesler, "Resource development and experiments in automatic sa broadcast news transcription," in *SLTU-Workshop on Spoken Language Technologies for Under-Resourced Languages*, Cape Town, South Africa, May 2012, pp. 102–106.
- [20] R. Eiselen and M. Puttkammer, "Developing text resources for ten south african languages," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation* (*LREC'14*). Reykjavik, Iceland: European Language Resources Association (ELRA), May 2014.
- [21] M. J. F. Gales, K. M. Knill, A. Ragni, and S. P. Rath, "Speech recognition and keyword spotting for low resource languages: Babel project research at cued," *Spoken Language Technologies for Under-Resourced Languages*, 2014.
- [22] M. Davel, E. Barnard, C. van Heerden, W. Hartmann, D. Karakos, R. Schwartz, and S. Tsakalidis, "Exploring minimal pronunciation modeling for low resource languages," in *Accepted to Inter*speech 2015, 2015.

# Introducing XGL - a lexicalised probabilistic graphical lemmatiser for isiXhosa

Lulamile Mzamo, Albert (ASJ) Helberg Faculty of Engineering North-West University Potchefstroom, South Africa lula mzamo@yahoo.co.uk, albert.helberg@nwu.ac.za

Abstract—In this paper, a lexicalized probabilistic graphical lemmatiser for isiXhosa, XGL, is presented. An overview of isiXhosa lemmatisation issues is given, followed by a discussion on previous work in automated lemmatisation for isiXhosa. The paper continues to motivate for a machine learning lemmatiser for isiXhosa. IsiXhosa data used to train the lemmatiser is analyzed and the best features are identified from the analysis. The inner workings of XGL are detailed and evaluation results presented. XGL is shown to have achieved accuracy rates of 83.19% on a gold standard of word-lemma pairs, thereby outperforming similar lemmatisers such as LemmaGen's 80.6% and 73.13% from the CST lemmatiser when trained with 35000 word-lemma pairs.

Keywords—Natural Language Processing; Machine Learning; Lemmatisation; IsiXhosa

#### I. INTRODUCTION

Human language resources and applications currently available in South Africa are of a very basic nature. According to Groenewald [1] this can be attributed to the dependence on Human Language Technology (HLT) expert knowledge, scarcity of data resources, lack of market demand for the African languages, and how the particular language relates to other more resourced languages. Lemmatisation is one of the basic tools in natural language processing (NLP). The work detailed in this paper is the development of a lemmatiser for one such language, namely isiXhosa.

IsiXhosa is one of the South African official languages belonging to the Bantu language family which are classified as "resource scarce languages" [1]. Although some work has been done in computational linguistic tools for isiXhosa it is of a limited nature [2]. IsiXhosa is the second largest language in South Africa with 8.1 million mother-tongue speakers (16% of the South African population), second only to isiZulu [3].

IsiXhosa is closely related to languages such as isiZulu, siSwati and isiNdebele and therefore work done in it could easily be bootstrapped to these languages as has been shown in [4].

Sonja (SE) Bosch Department of African Languages UNISA Pretoria, South Africa boschse@unisa.ac.za

#### II. LEMMATISATION FOR ISIXHOSA

#### A. Lemmatisation

Lemmatisation is "concerned with finding the lemma of a set of inflected word forms or with generally assigning lemmas to inflected word forms" [5].

In natural language processing, lemmatisation is looked at in terms of inflection as "a normalisation step on textual data, where all inflected forms of a lexical word are reduced to its common headword, the lemma" [6]. Jurafsky and Martin [7] explain this by stating that in the context of natural language processing a lemma represents a set of lexical forms with the same stem, the same major part-of-speech and the same wordsense.

A process similar to lemmatisation is stemming. For a particular paradigm, a stemmer simply finds the common substring among the paradigm word forms. The lemmatiser, in contrast, maintains the meaning. An example is that the lemma for "*better*", and "*best*" is "good" [8].

#### B. A lemma in isiXhosa

IsiXhosa is an agglutinating and polysynthetic language in that it has many morphemes per word [9], [10]. It is also fusional/inflectional because morpheme boundaries are fused and difficult to distinguish [9].

IsiXhosa words are composed of a root, prefixes, suffixes and circumfixes that attach to the root [11] which is the meaning carrying constituent of the word. A circumfix is the "simultaneous affixation of a prefix and suffix to a root or a stem to express a single meaning" [9]. An example of a circumfix in isiXhosa is the combination "a...nga" in isiXhosa negation, e.g. akahambanga (he/she did not go).

Each of the affixes (i.e. prefixes, suffixes or circumfixes) is made up of one or more morphemes. A morpheme is the smallest meaning bearing component of a word [9]. Morphemes follow one another in an order prescribed for each word type [12]. In isiXhosa, most roots are however bound morphemes which are not independently meaningful [9]. Stems are word roots suffixed with a termination vowel [12], and hence the use of stems as the appropriate lemma for isiXhosa.

#### C. Automated Lemmatisation for isiXhosa

One of the earliest reports on automated morphological analysis of South African languages is that of Theron and Cloete [13] on the automatic acquisition of a Directed Acyclic Graph (DAG) to model the two-level rules for morphological analysers and generators. The algorithm was tested on English adjectives, isiXhosa noun locatives and Afrikaans noun plurals. The algorithm was implemented for Afrikaans lemmatisation and achieved 5-fold validation accuracy of 93% for Afrikaans noun plurals [14].

The next lemmatisation work on isiXhosa was a supplement to spellchecking [15]. The primary objective was to identify lemmas so that inflection could then be applied to increase the lexicon of the spellchecker. This increased the lexical recall of the spellchecker from 78.82% to 92.52%.

The last lemmatisation work for isiXhosa is that used to generate the data for the study reported on in [16]. The exercise reported a rule-based lemmatiser accuracy rate of 79.82%.

#### **III. MACHINE LEARNING**

As stated before, the last work done in lemmatisation for isiXhosa with some success, is reported in [16]. That work was based on linguistic rules. The work presented in this paper took motivation from [17] to establish if it is possible to improve on the rule systems using a machine learning system. De Pauw and de Schryver [17] had shown that a machine learning lemmatiser for Kiswahili could be trained with the output of a rule-based lemmatiser, and perform better than the same rule-based lemmatiser, hence the development of XGL, the lexicalised probabilistic graphical lemmatiser for isiXhosa.

#### A. Machine Learning

Machine learning based systems learn the model of a task from experience. This model is then used in executing the task. A characteristic of machine learning is improvement in performance with more experience [18].

XGL is therefore expected to learn from word-lemma pairs (experience), with respect to lemmatisation (the task), and accuracy of lemmatisation (performance measure), if its accuracy (performance) as measured, improved with an increase in the number of word-lemma pairs (experience).

#### B. Machine Learning lemmatisers

The prevalent machine learning lemmatisation techniques primarily follow two technical streams, namely memory based learning/classifiers (MBLs) and Ripple-Down-Rules (RDRs).

Plisson et al. [19] introduced the Induced Ripple-Down-Rules (RDR) approach to word lemmatisation. RDR was originally used for rule-based systems and resembles "if-thenelse" statements with the most general rules appearing first and exceptions branching from them. In essence an RDR is a hierarchy of rules where the one level contains the rules and the following level contains exceptions to each rule, and so on. The work was done on Slovene and achieved accuracies of 77% at the time. Plisson et al. [20] modified the algorithm to handle exceptions better by recording words covered by a rule under that rule, resulting in accuracies of 97.2% in Slovene. Jongejan and Dalianis [21] presented a lemmatiser (CST Lemmatiser) that works with affixes, because languages like Dutch can include prefixing in addition to suffixing. CST is open source software. Jongejan and Dalianis [21] specifically state that the method used is not an obvious choice for agglutinating languages. This lemmatiser also used a hierarchy of rules. CST has been used in a number of other studies as well [22]–[24], achieving good results. Juršič et al. [25] presented an even more enhanced Ripple-Down-Rules lemmatiser called LemmaGen, tested on twelve languages. LemmaGen is also open source software.

Van den Bosch and Daelemans [26] introduced Memory based learning in a tool named the Tilburg Memory-Based Learner (TiMBL) for Dutch. TiMBL was successfully used by Groenewald [27] in the development of the Afrikaans lemmatiser LiA ("Lemma-identifiseerder vir Afrikaans", "Lemmatiser for Afrikaans"). LiA achieved accuracy rates of 91% for Afrikaans, but rates of 64.05% for Setswana [1].

The features that are selected affect performance in both RDR and MBLs as shown by the improvement achieved in lemmatisation for Afrikaans [27] and that shown by the reduction in lemmatisation accuracy for Italian for the suffixal LemmaGen lemmatiser [25].

Hybrid lemmatisation techniques show the best results, especially for highly inflecting languages as witnessed in the Lemmald lemmatisation for Icelandic [24].

For all three techniques, incorporating POS tags improves results substantially as witnessed in Lemmald [24] and also the work reported in [28].

#### C. Choice of lemmatisers for comparison

To place XGL among other lemmatisers, two of the better lemmatisers were chosen for comparison. The lemmatisers had to be publicly available and had to have been used for highly inflectional languages like isiXhosa.

The CST lemmatiser [21] was chosen because it is a good benchmark that has been used extensively, and is freely available.

The LemmaGen [25] lemmatiser was chosen because it implements the Ripple-Down-Rules algorithm, an algorithm that shows promise with highly inflected languages.

#### IV. XGL: THE LEXICALISED PROBABILISTIC GRAPHICAL LEMMATISER FOR ISIXHOSA

This section details the development of XGL. As a synthetic language isiXhosa uses prefixes, suffixes and circumfixes to produce more words from a stem. To reduce these words back to lemmas XGL uses transformation classes. These classes reflect the transformation between the full word form and a lemma. An example of a transformation class is shown below:

Ekuqinisekiseni => qina. : Leku>Risekiseni>a

#### A. Transformation Classes

The "L" in the transformation class specifies the start of the prefix transformation and the "R" specifies the start of the suffix transformation. The symbol ">" separates what should be removed, to the left of the symbol, and what should be inserted, to the right of the symbol.

In the above example, to transform "*ekuqinisekiseni*" to "*qina*", the prefix "*eku*" must be replaced with nothing hence the prefix transformation "*Leku*>". To the right of the word the suffix "*isekiseni*" must be replaced by "a", hence suffix transformation "*Risekiseni*>a". Combining the two transformations makes the word transformation class "*Leku*>*Risekiseni*>a".

#### B. Data and Feature selection

The data used in the study was the isiXhosa NCHLT Annotated Text Corpora<sup>1</sup>, a product of the NCHLT Project on Text Resources conducted by the North-West University's Centre for Text Technology (CTEXT) and the Republic of South Africa's Department of Arts and Culture [16]. Eiselen and Puttkammer [16] state that most of the data for the corpus was sourced from the South African government websites with the rest coming from news articles, scientific articles, magazine articles and prose. The data was generated using rules from a study conducted by Bosch et al. [29].

The corpus consisted of two lemma annotated files, a 50000 word training corpus of word-lemma pairs generated using a finite state lemmatiser [29] and a 5000 word gold-standard testing corpus of word-lemma pairs initially generated using a finite state lemmatiser [29] but quality assured by linguistic experts.

Data exploration was done on the training corpus. This confirmed that isiXhosa is a prefixal language with 3020 prefixes identified active in 84.8% of the data. 64% of the data was characterised by 2419 unique prefixes. There were 311 suffixes overall, active in 21% of the data. Only 0.24% of the data was characterised by suffixes only. Some of the prefixes and suffixes combined into 2504 circumfixes that covered 20.7% of the data. However, there are 5131 transformation classes overall, covering the data. Of that 14.9% of data lemmatised to the full word forms. No analysis was done on the testing gold standard.

It was also identified that 50% of the affixes tended to cover 95% of the data covered. An example of a data coverage graph is shown in Fig. 1.

<sup>1</sup>The data is available via the following website link: http://rma.nwu.ac.za/index.php/resource-

catalogue/isixhosa-nchlt-annotated-text-corpora.html



Fig. 1. Prefix coverage in prefix only data

The coverage of the affixes is shown in Table 1.

TABLE I. AFFIX COUNTS, AND THEIR MAXIMUM DATA COVERAGE

Affix	Count	Maximum Coverage
Prefixes	3020	84.8%
Suffixes	311	21%
Circumfixes	2504	20.7%
All Affix Classes	5131	85.1 %

Because prefixes cover so much of the data, performance of searching for the lemmatisation class could be drastically improved by using the prefix as the primary search index, followed by the suffix.

An analysis of the word length relative to the combined length of the affixes presented a Pearson correlation of 68.1%.



Fig. 2. Bubble plot of Affix length by word length

A bubble plot of the relationship between word length and affix length is shown in Fig. 2. The correlation showed that one could also use length to decide on the appropriate set of affixes for a word.

The coverage of the affixes in combination with the correlation of the word length to the combined affix length

gave the indication that a prefix-suffix-word length feature set would give good class discrimination when finding the right lemmatisation class for a word.

#### C. How XGL works

Fig. 3 shows the workflow of XGL. The centre of XGL is, however, the model.



#### Fig. 3. XGL workflow

XGL's lemmatisation model is made up of two components i.e. the lexicon and the conversion class tree.

A sample of the lexicon is shown below. The lexicon is indexed by the full word form and references lemmas encountered during training with counts of how many times such a lemma was encountered for that word, as shown below:

```
'aliphelise': { 'phelisa': 1},
'aliqingqelwa': { 'qingqa': 1},
'aliyi': { 'ya': 1},
'alo': { 'lo': 2},
```

A sample leaf of the conversion class tree is show below:

```
('be', 'isa'):
    {'CLASSES':
        {'Lbe>Risa>a': {'Count': 1,
            'Stats': {'Mean':
            12.0,'Std': 0.0}},
        'Lbe>Risa>o': {'Count': 1,
            'Lbe>Risa>o': {'Count': 1,
            'Stats': { 'Mean': 11.0,
            'Stat': 0.0}
    }
},
```

Each leaf is indexed by the class affixes, and contains conversion classes. Each class contains statistics on the length of the word. This is used to model the probability of the word belonging to the class, p(*class*|word).

The leaves in the class tree are arranged in a hierarchy by the affix tree as shown below:

The Lemmatisation part of XGL works is shown in Fig. 4.

The lexicon is the first port of call when looking for a word lemma. When that fails the class tree is used. Searching the class tree for matching classes involves enumerating all the classes whose prefix-suffix combinations respectively match the beginning and end of the test word. The probability of the test word belonging to a class p(class|word) is then calculated and the highest one selected. The lemmatiser uses a probability threshold of 0.975 for the probability of a word belonging to a class to decide if the word should be lemmatised or returned as is. This threshold showed the best results when tuning the XGL lemmatiser.



Fig. 4. Word lemmatisation workflow

#### V. EVALUATION

As mentioned before, XGL was evaluated against the CST and LemmaGen lemmatisers.

#### A. Experiment setup

The training corpus of the isiXhosa NCHLT Annotated Text Corpora was used for training the lemmatisers.

Ten-Fold training was done, and testing was conducted on the testing corpus.

#### B. Results

Fig. 5 shows 10-Fold average accuracy trends for the three lemmatisers for each training sample size.



Fig. 5. Lemmatisation accuracy results by training set size

The maximum average accuracy rates were attained at the training set size of 35000. They were 83.19% for XGL, followed by LemmaGen at 80.6% and the CST lemmatiser at 73.13%. As evident in Fig. 5, the isiXhosa Graphical Lemmatiser had average accuracy trends that outperformed the other two lemmatisers when tested on the isiXhosa NCHLT Annotated Text Corpora.

To verify the hypothesis that XGL performed significantly better than the LemmaGen and CST lemmatisers the paired Wilcoxon signed-rank test [30] was used. The Wilcoxon signed-rank test is recommended by Dems'ar [31] who found that the widely used t-test was an inappropriate and statistically unsafe comparison test for classifiers. The Wilcoxon calculations were done to compare XGL to the other two lemmatisers with resultant p-values of at most 0.005 across the range of training set sizes. This is an order of magnitudes smaller than the threshold of 0.05.

This indeed showed that XGL performed statistically significantly better than LemmaGen and the CST lemmatiser in this specific corpus.

#### VI. CONCLUSIONS

A lexicalised probabilistic graphical lemmatiser for isiXhosa, XGL, was presented. On the isiXhosa lemmatisation corpus of the NCHLT Annotated Text corpora, XGL performed significantly better in accuracy than both the CST and LemmaGen lemmatisers, two lemmatisers that have shown significant use and performance on highly inflectional languages. An average accuracy of 83.19% on a training data set size of 35000 word-lemma pairs and the increase in accuracy with the size of training data suggests that more data would improve accuracy even further.

This study also confirms that one can use data generated using a rule-based system and even achieve better performance than the rule-based system, as initially shown by De Pauw and de Schryver [17]. The training data used in this study was generated using a rule-based system and was evaluated to an accuracy of 79.82% which is lower than the 83.19% achieved by XGL using the same data.

This work was restricted to one language (isiXhosa) using a small training set and primarily government data. Future work could include testing XGL's performance in other closely related languages, as well as testing XGL on a more balanced and considerably larger data set.

#### Acknowledgment

We want to thank the South African Language Resource Management Agency (http://rma.nwu.ac.za/index.php/) for providing a central source of data and resources for Natural Language Processing work.

#### References

- H. J. Groenewald, "Using Technology Transfer to Advance Automatic Lemmatisation for Setswana," in *Proceedings of the EACL 2009 Workshop on Language Technologies for African Languages – AfLaT 2009, pages 32–37, Athens, Greece, 31 March* 2009., 2009, pp. 32–37.
- [2] A. Sharma Grover, G. B. van Huyssteen, and M. Pretorius, "HLT profile of the official South African languages.," in 2nd AFLaT workshop at the Seventh International Conference on Language Resources and Evaluation (LREC) 2010, 2010, pp. 3–7.
- [3] Statistics South Africa., "Census 2011: Census in brief," 2012.
- [4] S. Bosch, L. Pretorius, and A. Fleisch, "Experimental Bootstrapping of Morphological Analysers for Nguni Languages," vol. 17, no. 2, pp. 66–88, 2008.
- [5] S. R. Spiegler, "Machine Learning for the analysis of morphologically complex languages," PhD Thesis, University of Bristol, 2011.
- [6] T. Erjavec and S. Dzeroski, "Machine learning of morphosyntactic structure: Lemmatizing unknown Slovene words," *Appl. Artif. Intell.*, vol. 18, no. 1, pp. 17–41, 2004.
- [7] D. Jurafsky and J. H. Martin, Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Upper Saddle River, New Jersey: Pearson Prentice Hall, 2000.
- [8] W. Daelemans, H. J. Groenewald, and G. B. Van Huyssteen, "Prototype-based Active Learning for Lemmatization," *Proc. RANLP* '2009 Int. Conf. Recent Adv. Nat. Lang. Process., pp. 65–70, 2009.
- [9] I. M. Kosch, Topics in Morphology in the African Language Context. Pretoria: Unisa Press, 2006.

- [10] E. M. Bender, Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Morphology and Syntax. Morgan & Claypool, 2013.
- [11] H. W. Pahl, *IsiXhosa*. King Williams Town: Educum Publishers, 1982.
- [12] J. A. Louw, R. Finlayson, and S. C. Satyo, *Xhosa Guide 3 for XHA100-F.* Pretoria: University of South Africa, 1984.
- [13] P. Theron and I. Cloete, "Automatic acquisition of two-level morphological rules," in *Proceedings of the fifth conference on Applied Natural Langauge Processing*, 1997, pp. 103–110.
- [14] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, Third edition, Prentice H. London: Prentice Hall Press, 2014.
- [15] J. Jones, K. Podile, and M. Puttkammer, "Challenges relating to standardisation in the development of an isiXhosa spelling checker," *South African J. African Lang.*, vol. 25, no. 1, pp. 1–10, 2005.
- [16] R. Eiselen and M. J. Puttkammer, "Developing Text Resources for Ten South African Languages," in *Proceedings of the Ninth International Conference on Language Resources and Evaluation* (*LREC'14*), 2014, pp. 3698–3703.
- [17] G. De Pauw and G. de Schryver, "African Language Technology: The Data-Driven Perspective," in Second Colloquium on Lesser Used Languages and Computer Linguistics, 2009, pp. 79–96.
- [18] T. M. Mitchell, Machine Learning. WCB/McGraw-Hill, 1997.
- [19] J. Plisson, N. Lavrac, and D. Mladenic, "A Rule based Approach to Word Lemmatization," in *Proceidings of the 7th International Multi-Conference Information Society JS 2004*, 2004, pp. 83–86.
- [20] J. Plisson, N. Lavrac, D. Mladenić, and T. Erjavec, "Ripple Down Rule learning for automated word lemmatisation," *AI Commun.*, vol. 21, pp. 15–26, 2008.
- [21] B. Jongejan and H. Dalianis, "Automatic training of lemmatization rules that handle morphological changes in pre-, in- and suffixes alike.," in *Proc. 47th Annual Meeting of the ACL and the 4th IJCNLP of the AFNLP. Suntec*, 2009, pp. 145–153.
- [22] S. Saraswathi and T. V. Geetha, "Comparison of performance of enhanced morpheme-based language model with different word-

based language models for improving the performance of Tamil speech recognition system," *ACM Trans. Asian Lang. Inf. Process.*, vol. 6, no. 3, p. Article number 9, Nov. 2007.

- [23] Z. Agic, N. Ljubesic, and Danijela Merkler, "Lemmatisation and Morphosyntactic Tagging of Croatian and Serbian," in *Proceedings* of the 4th Biennial International Workshop on Balto-Slavic Natural Language Processing, 2013, pp. 48–57.
- [24] A. K. Ingason, S. Helgadóttir, H. Loftsson, and E. Rögnvaldsson, "A mixed method lemmatization algorithm using a hierarchy of linguistic identities (HOLI)," in Advances in Natural Language Processing, 6th International Conference on NLP, GoTAL 2008, Proceedings, 2008, pp. 205–216.
- [25] M. Juršič, I. Mozetič, T. Erjavec, and N. Lavrač, "LemmaGen: Multilingual Lemmatisation with Induced Ripple-Down Rules," J. Univers. Comput. Sci., vol. 16, no. 9, pp. 1190–1241, 2010.
- [26] A. Van den Bosch and W. Daelemans, "Memory-based Morphological Analysis," in *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, 2009, pp. 285–292.
- [27] H. J. Groenewald, "Educating Lia: The development of a linguistically accurate memory-based lemmatiser for Afrikaans," in *Intelligent Information Processing III*, vol. 228, Z. Shi, K. Shimohara, and D. Feng, Eds. Boston: Springer US, 2007, pp. 431– 440.
- [28] A. Gesmundo and T. Samardžić, "Lemmatisation as a tagging task," in Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers, 2012, vol. 2, pp. 368– 372.
- [29] S. Bosch, J. Jones, L. Pretorius, and Winston Anderson, "Resource Development for South African Bantu Languages: Computational Morphological Analysers and Machine - Readable Lexicons," in Proceedings on the Workshop on Networking the Development of Language Resources for African Languages. 5th International Conference on Language Resources and Evaluation, 2006, pp. 38– 43.
- [30] F. Wilcoxon, "Individual Comparisons by Ranking Methods," *Biometrics Bull.*, vol. 1, no. 6, pp. 80–83, 1945.
- [31] J. Dems'ar, "Statistical Comparisons of Classifiers over Multiple Data Sets," J. Mach. Learn. Res., vol. 7, pp. 1–30, 2006.

## BEETLE - A Modular Electronics Family for Robotics

John Dickens and Hein Swart CSIR Material Science and Manufacturing Pretoria, South Africa Email: jdickens@csir.co.za

Abstract-Mobile robotics has a wide range of applications, resulting in a diverse array of designs including a variety of sensors and manipulators. The task of integrating the variety of components that make up a typical robotic system takes significant effort and resources. In order to be able to rapidly react to changing market demands for automation it is necessary to have a system that allows simple integration of various components while retaining enough versatility to address a variety of applications. A family of modular electronic elements is proposed to address this need. The Beautiful Embedded Electronic Logic Element (BEETLE) family of boards is designed to be compact, low cost, robust, reusable and easy to program. This allows the boards to be used in a wide variety of applications including space and cost constrained systems. The family uses a combination of the Robotic Operating System (ROS) and Arduino compatibility which allows easy integration with robotic systems and simple modification of functionality as needed.

Index Terms-modular, versatile, electronics, robotics, ROS, Arduino

#### I. INTRODUCTION

The field of mobile robotics has a vast array of applications ranging from surveillance and security through search and rescue to bomb disposal and underground mine inspection. The wide variety of applications results in a large number of robot configurations. Like any industry, there is increasing competition to reduce costs, increase quality and reduce development times. Modularity is often mentioned as a means of balancing the apparently opposing requirements of standardisation for mass production with flexibility and customisation [1, 2]. Mass production allows the reduction of costs through economies of scale while flexibility and customisation enables companies to rapidly respond to customer requirements. The widespread use of modular design in the computer industry has often been cited as a critical element in the rapid development of computing technology [3, 4]. The basic drivers of modularity are the creation of variety, the reduction of complexity and the reuse of similarity. All of these drivers are present in a typical robotic system, as discussed there is a wide variety of configurations but there are also significant similarities with many sensor and actuator systems being common to many robot configurations. There are a number of benefits that can be gained from modularity, some of these potential benefits include:

- Economies of scale
- Increased feasibility of product/component changes
- Increased product variety

- Reduced lead time
- Decoupling of tasks
- The ease of product upgrade, maintenance, repair, and disposal

A system can only be called modular by looking at the system as a whole. What may appear to be a component of a system could in fact be a module of a larger system, the essential characteristics of a modular methodology are the following [1]:

- The use of a finite number of modules/components to meet an infinite number of configurations.
- The modules keep as much independence as possible.
- General usability.
- Seamless interfacing between elements.
- Modules incorporate self-contained functionality

The concept of modularity in robotics is widely published but the focus has mostly been on the design of modular robot systems[5, 6] or actuators [7]. A modular electronic system for robotics has been developed that makes use of an FPGA based distributed control boards [8]. The use of FPGAs for distributed control provides advantages of hard real-time control ability, processing power and versatility. However the high cost of FPGAs such as the Spartan 6 used by Pierce and Cheng[8] make them a costly option for a robotic system requiring many simple modules.

There are a number of commercial systems that make use of a modular methodology for embedded electronics but while they have many applications they all have disadvantages that the authors believe make them unsuitable for use in many mobile robotic systems. One example of an embedded electronic system that can be seen as using a modular approach is the Arduino with its shields [9]. For Arduino boards and shields there is a standard interface and interaction between the board and shield, there are a number of boards that are compatible with a number of different shields and their functionality is self-contained in the board or shield. The main disadvantage of Arduino is that the boards are designed for prototyping and not for a field application. Arduinos are not sufficiently robust to be used in a field application, experince has shown that the connectors (especially the USB) have a tendency to vibrate loose in high vibration environments. An even more modular system for embedded electronics development is the Processor Independent Embedded Platform (PIEP) [10] which is a motherboards and a number of processor and peripheral boards that plug in to the main motherboard. Like Arduino the PIEP system is designed for prototyping and appears suffer from the same issues as Arduino. An example of an industrial modular system would the CompactRIO from National Instruments [11], the system consists of a reconfigurable chassis, various IO modules and a real-time controller. The CompactRIO has the advantages of providing a robust industrialised system with mature hardware and software, however there are a number of disadvantages limiting its use in general robotics applications. The disadvantages are cost, size, power consumption and propriety interfacing and software. For example the CompactRIO Controls and Mechatronics Bundle that consists a controller, chassis and analogue and digital input/output modules costs in the order of R150 000, requires a 40 W power supply and starts at 180 x 93 x 87 mm in size. The issue of cost obviously restricts the use of the system in many applications and the large size and high power consumption makes the CompactRIO unsuitable for small, battery operated systems. The issue of propriety interfacing and software makes integrating the CompactRIO with other systems difficult.

This paper describes the design of a family of modular embedded electronics boards for robotic (and other) applications. Section II provides an overview of the Beautiful Embedded Electronic Logic Element (BEETLE) family of boards, this is followed by a discussion of the interface and interaction standards for the modules. A discussion of the design of the hardware follows with a specific focus on the design of the Digital IO and Motor Controller boards.

#### **II. SYSTEM OVERVIEW**

The BEETLE family of boards is designed to meet a number of requirements for general electronics within the Mechatronics and Micro-Manufacturing (MMM) group. Within the group, there are a wide variety of projects that include robotics but are not limited to it. It would be advantageous if the boards could be used in applications other than just robotics. The important requirements for a system of general electronic modules are:

- Low cost
- Reusable
- Compact
- Expandable new functionality easily added
- Easy to use
- Building block for robotic and other devices
- On board processing (basic)
- Robust

The majority of the robotic systems developed by the MMM group use the robotic operating system (ROS). ROS is a meta operating system (or middleware) with the primary goal of simplifying the task of creating complex and robust robot behaviour [12]. ROS features an inter-process and inter-machine communication system, a package management system as well as tools and libraries for obtaining, building, writing, and running code across different computers. One of the core design features of ROS is to be modular and peer-to-peer

with one process per computational element and a well defined inter-process communication system. For robotic applications the boards are designed to use rosserial which allows the board to connect to a host running ROS and transparently interact with the system through ROS's publish and subscribe system. To ensure that the boards are easy to program for other applications the BEETLE family is Arduino compatible.

The electronic sub-system of a robot will be laid out in a star configuration with a host computer connected to multiple BEETLE boards as slaves. The availability of small, low cost, low power ARM based single-board computers makes this configuration appropriate for even small robots, however the embedded processing on each module means that it is also possible to use a module (or modules) without a host.

#### A. Interface and Interaction

The concept of modularity requires the creation of variety through combination. Modules can only be interchanged to create combinations if the interfaces and interactions are standardised. Interfaces are the boundaries linking modules together and interactions describe the input/output relations between modules.

For the BEETLE boards there are two levels of interface and a single interaction. The interfaces are the mechanical and electronic interfaces, the mechanical interface is the mounting which is on a grid of 50 mm x 50 mm and the electronic is a universal serial bus (USB) interface that provides the data and power for the processor on the boards. The software interaction is provided by a standardisation on a 'rosserial' communication protocol to a host machine running ROS.

The advantage of using rosserial is that it allows the boards to become transparent to the rest of the system. All that the host system needs is the port on which the board is connected and to run a rosserial host node. The rosserial code on the board will publish information to the host and subscribe to commands from the host and the host system does not need to know how the board deals with the information. For example a BEETLE motor control board connected to a motor and encoder could be plugged into a host running ROS and it would transparently identify its capabilities. It would advertise sensor values and subscribe to control topics, such as advertising a joint position and subscribing to a joint command topic. Additionally through the use of ROS parameters the operation of the board can be modified at run time. From the host software perspective the operation of the board is abstracted to an interaction through the ROS topic protocol. This allows the boards to be reusable and expandable with new functionality while still retaining a well defined interaction protocol necessary to be truly modular.

#### B. BEETLE Hardware

The BEETLE family of electronic boards was designed around a few principals, as discussed above. There is a core module that is central to all designs in the family. The flexibility of interfaces is vital to provide a wide variety of applications in the field, but a standardised software interface



Fig. 1. The core BEETLE module functionality

is critical to retain a common programming interface and useful modularity across all boards. Fig. 1 shows the layout of the core electronics. The BEETLE family of boards are intended for robotics and industrial applications, which may experience significantly more vibrations, electrical interference and spikes when compared to hobby or lab level applications. All connectors and interfaces are designed to cater for this environment. This includes latching pluggable connectors, using cage clamps instead of screw terminals where available and having protection circuitry on all input and output connections on the board.

Central to the electronics is the ATSAMD21G18A ARM Cortex-M0+ based microcontroller from Atmel (#1 in Fig. 1) with 256 KB Flash, 32 KB SRAM, 48 MHz internal clock, three 16-bit timer/counters, DMA, 6 serial communications interfaces (SERCOMs), USB host and device, 14-channel 12bit analogue-to-digital converter (ADC) and 10-bit digitial-toanalogue converter (DAC). This gives the BEETLE boards the processing power and a lot of flexibility to be used for a wide range of applications. Standardised communication is central to the BEETLE philosophy, thus USB communication was selected as the main interface to all BEETLE boards. USB was selected as it is a common interface on higher level computing platforms plus it supplies power and communication in a single connection. This enables the BEETLE boards to operate with minimal external peripherals (such as extra power supplies) and without special adapters that enable communication. From experience, the drawback with USB is the clearance required for connectors, the difficulty in creating custom length cables and routing standard cables through panels inside an application and the non-latching nature of the connectors that have a tendency to vibrate loose over time. To solve these challenges, two USB connection interfaces are supplied on each BEETLE board. First there is a standard micro-USB port (#2 in Fig. 1) for bench development and standalone applications where a standard cable is sufficient. Secondly there is a 2 mm pitch latching header from the Molex MicroClasp wire-to-board range (#3 in Fig. 1) that enables easily prepared custom cables, that are reliable in a high vibration environment, to be used. Both connectors terminate into an IEC 61000-4-2 level 4 voltage protection (15 kV air and 8 kV contact discharge) circuit (#4 in Fig. 1) for both the USB power and data lines. A debug header for the Atmel ICE Debugger (#5 in Fig. 1) is included that enables bootloader programming as well as full debugging capabilities. Finally a 32 pin header (#6 in Fig. 1) brings out all available IO pins and the 3.3 V and 5 V power rails. The core also includes the following capabilities: three LEDs (#7 in Fig. 1) for Status (Blue), USB Transmit (Green) and USB Receive (Red); 3.3 V regulator from USB power (#8 in Fig. 1); 32.768 kHz real-time clock crystal (#9 in Fig. 1); and solder jumper for 3.3 V supply as analogue reference voltage (#10 in Fig. 1). This core board is the starting point for all designs and is shown in Fig. 2.

The goal of the BEETLE family tree is to divide all possible applications into a set number of versatile electronic boards that have a standardized communication and software interface. This family is divided into seven distinct groups: digital signal, analogue signal, motor control, power, sensor, communication and user interface as shown in Fig. 3.

- The digital signals group focuses on reporting and changing discrete signals. Boards that falls into this category is a general IO board that reads a set number of discrete inputs and sets a number of discrete outputs for low current applications and a relay board enabling the isolated the switching of high power signals.
- The analogue signal group focusses on constantly varying



Fig. 2. The BEETLE core board

signals with a general ADC/DAC combination board as well as a low channel-count, high resolution ADC board.

- The motor controller group of boards implement all the functions associated with controlling a DC motor. This includes interfaces to speed and position sensors (incremental and absolute encoders, potentiometers, etc.) as well as monitoring the condition of the motor and supply (supply voltage, current, temperature, etc.). A series of higher power boards makes up the collection of board within this group.
- The power group relates to all things required to manage the supply of power to the application. The first board is a smart battery management board that manages the charge and discharge cycles of a battery pack as well as monitoring and controlling the health (temperature, cell balancing, available energy levels, etc.) of the pack. Multiple batteries can be managed as a whole, while each one is looking after itself, disconnecting from the bigger system to "save" itself in the case of a problem. An advantage of this is that high power switching (such as start-up) can happen digitally, thus controllably, instead of a single surge as with a traditional mechanical switch. A predecessor of this board was implemented with great success in previous applications. A second power management board monitors power usage in non-battery applications as well as provides a set of configurable and switch-able voltages for the application and manages start-up sequences.
- The sensor group focusses on measuring physical quantities and environmental changes such as movement (9 degree-of-freedom (DOF) inertial measurement unit (IMU) board), gases (gas sensor board), displacement and flow (ultrasonic distance board) and location (Global Differential GPS board)
- The communication group provides extra functionality to the family to enable integration of unique peripherals to the standard software interface of the BEETLE family.

This includes a high port count USB hub with the same Molex MicroClasp wire-to-board range connectors on the separate boards to further enhance the easy wiring and latching of the USB cables. This also includes power management of each port to ensure a reliable communication system. The next board in this group is a serial communication gateway board that will enable the interface to standard serial interfaces (such as RS232, RS485, CAN, SPI, I2C, etc.) at their correct voltage levels. Finally a set of wireless gateway boards that provides wireless communication links (such as, 6LoWPAN, Bluetooth, WiFi, ZigBee) with other platforms on various ISM bands (433 *MHz*, 868 *MHz*, 2.4 *GHz*, 5 *GHz*).

• The user interface board is designed for simple user interface with the larger system. The initial interface board will be a simple user interface with basic pushbutton inputs from the user and a character LCD for user feedback. This may be extended later but it is envisioned that more complex user interfaces would be graphical user interfaces running on a host computer.

The first two boards developed within the BEETLE family are the Digital IO board shown in Fig. 4 and the 50 A Brushed Motor Controller board (Fig. 5). The Digital IO board has 8 digital inputs, 8 digital outputs and two adjustable linear regulators to supply peripherals with power. All the inputs, outputs and power interfaces are capable of 0-30 V DC range and are protected from voltage spikes outside this range. The connectors used are 2.5 mm pitch cage clamp connectors that require no tools to insert or remove the wires.

 TABLE I

 Summary of the Digital IO board specifications

Parameter	Value
Number of inputs	8
Number of outputs	8
Regulator input voltage	3-30 V (x2)
Regulated output voltage	1.5-28.5 V (x2)
Input low	0-1.5 V
Input high	4-30 V
Input impedance	156 $k\Omega$
Output type	Open drain
Max output current	100 mA

The 50 A brushed Motor Controller is capable of driving DC motors within a voltage range of 11-36 V DC (with spikes of up to 60 V) and current range of 0-50 A. The supply voltage, current drawn as well as motor and board temperature is constantly monitored and reported over the BEETLE USB communication link. Extra configurable interfaces are supplied with a 2.5 mm pitch cage clamp connector that enables discrete digital signals (for example QA, QB, I from a quadrature encoder) as well as serial communication signals (SPI, UART) and an analogue input (for example a potentiometer) to interface to peripherals normally associated with motor control. The high-power motor and all input signals are over-voltage protected and isolated (1000 V Isolation) from the core BEETLE electronics to minimize motor spikes from being injected into and affecting the rest of the system.



Fig. 3. BEETLE product tree

TABLE II Summary of the 50 A Motor Controller board specifications

Parameter	Value
Input voltage	11-36 V
Motor current	0-50 A
Quadrature encoder input	1
Analogue potentiometer input	1
Pulse width modulation frequency	$0-50 \ kHz$
Current measurement range	0-50 A
Current measurement resolution	25 mA
Voltage measurement range	0-48 V
Voltage measurement resolution	12 mV

#### III. CONCLUSION

The design of a family of modular electronics boards for robotics, and other, applications is discussed. The boards



Fig. 4. The BEETLE digital IO board



Fig. 5. The BEETLE 50 A brushed motor controller board

have a standardised processing core that is used throughout the family, this together with the common software interface makes the boards simple to interface and the functionality easy to modify. Each board is designed to provide a certain type of functionality. The specific functionality of the board can be easily customised in software running on the boards embedded processor. The boards are designed to be mechanically and electrically robust enough to be directly deployed in robotic and automation systems. A full family of boards has been envisioned and two have currently been produced. The Digital IO and 50 A Brushed Motor Control boards have been produced and lab tested and the process of deploying them in existing robotic systems is under-way.

#### IV. FUTURE WORK

Future work will include the development of the full family of BEETLE boards, as well as the deployment of the boards in existing and new robotic systems developed by the MMM group. Additionally the future version of the motor controller board will have support for reading from serial-output digital absolute encoders.

#### References

- Z. Bi and W. Zhang, "Modularity technology in manufacturing: taxonomy and issues," *The International Journal* of Advanced Manufacturing Technology, vol. 18, no. 5, pp. 381–390, 2001.
- [2] T. D. Miller and P. Elgard, "Defining modules, modularity and modularization," in *Proceedings of the 13th IPS research seminar, Fuglsoe*, (Fuglsoe, Denmark), 1998.
- [3] C. Y. Baldwin and K. B. Clark, *Design Rules: The power* of modularity, vol. 1. MIT Press, 2000.
- [4] R. Garud, A. Kumaraswamy, and R. Langlois, *Managing in the Modular Age: Architectures, Networks, and Organizations.* Wiley, 2009.
- [5] R. F. M. Garcia, A. Lyder, D. J. Christensen, and K. Stoy, "Reusable electronics and adaptable communication as implemented in the odin modular robot," in 2009 IEEE International Conference on Robotics and Automation, (Kobe, Japan), pp. 1152–1158, IEEE, May 2009.
- [6] M. Yim, D. G. Duff, and K. D. Roufas, "Polybot: a modular reconfigurable robot," in *Proceeding of the*

2000 IEEE International Conference on Robotics and Automation, (San Francisco, California), April 2000.

- [7] C. J. Paredis, H. B. Brown, and P. K. Khosla, "A rapidly deployable manipulator system," in *Proceedings of the* 1996 IEEE International Conference on Robotics and Automation., (Minneapolis, Minnesota), April 1996.
- [8] B. Pierce and G. Cheng, "Versatile modular electronics for rapid design and development of humanoid robotic subsystems," in 2014 IEEE/ASME International Conference on Advanced Intelligent Mechatronics, (Besanon, France), Jul 2014.
- [9] Arduino LLC, "Arduino introduction." Online, Aug 2015. http://www.arduino.cc/en/Guide/Introduction.
- [10] E3 Embedded Systems LLC, "Processor independent embedded platform." Online, Aug 2015. http://www.e3embedded.com/.
- [11] National Instruments Corporation, "Embedded systems for monitoring and control." Online, Aug 2015. http://www.ni.com/embedded-systems/.
- [12] "About ros." Online, Aug 2015. http://www.ros.org/about-ros/.

### Predicting vowel substitution in code-switched speech

Thipe I. Modipa<sup>1,2</sup> and Marelie H. Davel<sup>2,3</sup> <sup>1</sup>HLT Research Group, CSIR Meraka, South Africa. <sup>2</sup>Multilingual Speech Technologies, North-West University, Vanderbijlpark, South Africa. <sup>3</sup>CAIR, CSIR Meraka, South Africa.

tmodipa@csir.co.za, marelie.davel@gmail.com

Abstract—The accuracy of automatic speech recognition (ASR) systems typically degrades when encountering codeswitched speech. Some of this degradation is due to the unexpected pronunciation effects introduced when languages are mixed. Embedded (foreign) phonemes typically show more variation than phonemes from the matrix language: either approximating the embedded language pronunciation fairly closely, or realised as any of a set of phonemic counterparts from the matrix language. In this paper we describe a technique for predicting the phoneme substitutions that are expected to occur during code-switching, using non-acoustic features only. As case study we consider Sepedi/English code switching and analyse the different realisations of the English schwa. A code-switched speech corpus is used as input and vowel substitutions identified by auto-tagging this corpus based on acoustic characteristics. We first evaluate the accuracy of our auto-tagging process, before determining the predictability of our auto-tagged corpus, using non-acoustic features.

#### I. INTRODUCTION

Code switching tends to occur wherever speakers are exposed to more than one language. That is, multilingual speakers tend to mix languages naturally: embedding words or phrases from one language into speech primarily produced in a different language. As code switching introduces additional variability with regard to all aspects of speech – vocabulary, word usage and pronunciation – the presence of code-switched speech poses a challenge to automatic speech recognition (ASR) systems.

In an earlier study of Sepedi/English code-switched speech speech [1], an analysis of Sepedi radio broadcasts indicated that an unexpectedly high percentage (approximately 30%) of Sepedi utterances contained English words or phrases. It was also found that these words/phrases degrade ASR recognition significantly, resulting in an approximately 10% absolute degradation in performance. In the study, no special provision was made for these code-switched content, apart from including both English and Sepedi speech in the training data. Interestingly, by adding a 'Sepedi' pronunciation of English words by blindly performing grapheme-to-phoneme (G2P) prediction using Sepedi G2P models, some recognition improvement was observed by the authors.

This last observation was the primary motivation for the current study. In order to determine whether a more sophisticated approach to pronunciation modeling of the English words could improve results further, we first ask how predictable phoneme substitutions really are. Specifically we would like to determine whether vowel substitutions can be predicted based on non-acoustic features such as vowel context, word orthography or even speaker characteristics. We focus on vowels, as they exhibit significantly more variability in pronunciation than consonants, and specifically on the schwa, as its pronunciation tends to be the most unpredictable of all English vowels found in code-switched speech [2]. We approach this task by (a) first labeling a speech corpus with vowel tags based on an acoustic analysis, and then (b) determining how predictable these tags are using non-acoustic features.

The paper is structured as follows: Section II provides related background, specifically with regard to recognising code-switched speech, Sepedi/English code-switching and the 'Goodness of Pronunciation' score: a technique utilised in this study. Section III describes the approach we followed for the analysis and prediction of vowel substitution in codeswitched speech. Findings are presented in Section IV, before we conclude with a summary of our findings in Section V.

#### II. BACKGROUND

Code switching can be defined as the use of words or phrases from more than one language [3]. Such *code switching* can occur within a sentence, which is normally referred to as *intra-sentential* code switching. When the switching of language occurs between sentences the process is referred to as *inter-sentential* code switching [4]. During code switching, speakers may use words, phrases or sentences from one language (the *embedded language*) in combination with words or sentences in their primary or *matrix language* [3].

Code switching is a common phenomenon and brings with it challenges for ASR systems. These challenges can be overcome by building multilingual speech recognition systems consisting of multilingual pronunciation dictionaries, multilingual acoustic models, and multilingual language models. As an alternative, monolingual speech recognition systems can be run in parallel thereby switching from one system to the other [5], [6].

We study the development of combined, multilingual systems, and specifically the process to construct a multilingual phone set. The most popular approaches to the development of the phone set and phone-to-phone mappings, are as follows:

• Combining the phone set from multiple languages [7].

- Mapping the embedded phone set to the matrix phone set using IPA features directly. (IPA features classify sounds based on the phonetic characterisation of those speech sounds [5])
- Mapping highly confusable phones from the embedded to the matrix language based on a confusion matrix obtained from an existing ASR system.
- Merging language-dependent phone sets using hierarchical phone clustering algorithms and acoustic distance measures [7].
- Mapping phones between source and target sequences using probabilistic phone mapping [8]

In earlier work [1], [2], initial results with regard to the implications of Sepedi-English code switching for ASR systems were obtained on two custom-developed corpora. Specifically, in [2] a subset of an existing Sepedi corpus was selected and processed in order to isolate English events; and in [1] a new corpus was developed containing instances of code-switching as observed in radio broadcasts. This corpus, Sepedi Prompted Code-Switched (SPCS), was aimed to capture a speaker-specific pronunciation variability introduced during code switching.

The Goodness-of-Pronunciation (GOP) score was initially developed by Witt and Young in the context of phonelevel pronunciation assessment [9]. Defined as the durationnormalised log of the posterior probability (that the speaker uttered a specific phone, given the acoustic data), it is approximated by the difference in log likelihood of the target and best matching phone, divided by the number of frames in the segment, that is:

$$GOP(q) = \left| log \frac{(p(x|q))}{(p(x|q'))} \right| / NF(x)$$
(1)

where q is the target phone, x the observed data, NF(x) the number of frames observed and q' the model that matches the observed data best. In practice, the log likelihood scores are obtained directly from the ASR system, a hidden Markov model (HMM)-based one in our case, and q' identified during a free phone decode.

GOP was developed for phone-level analysis. In [10] a word-level version of GOP was defined, with two variants – frame-based and phone-based – depending on how duration normalisation is applied. In the same study, it was found that triphones provide more accurate results than monophones for word-level analysis. (Monophones are more typical of GOP scores used for phone-level pronunciation assessment). Word-level GOP, using triphones and frame-level normalisation, forms the basis for the analysis peformed here.

#### III. APPROACH

In code-switched speech, vowels are typically produced in one of two ways: the true embedded language pronunciation is produced or at least approximated fairly closely, or the target phone is substituted for a counterpart from the matrix language. As there is always the possibility of producing the true pronunciation, modeling both these possibilities requires that variants be introduced to the pronunciation lexicon: both 'embedded language' and 'matrix language' versions for each code-switched word are therefore required. Examples of such pronunciation variants are shown in Table I, using SAMPA notation<sup>1</sup>.

TABLE I EXAMPLES OF EMBEDDED AND MATRIX PRONUNCIATIONS

wor	ď	embedded language	matrix language
		(English)	(Sepedi)
pres	ssure	/p r\E S @ r\/	/prESa/
fifte	en	/ f @ f t i: n /	/fiftin/

Our goal is then to determine which features influence vowel substitutions when they *do* occur. That is, we would like to predict which vowel substitutions can be expected in the matrix pronunciation, specifically. Given the two examples in Table I, we therefore would like to predict that  $/@/ \rightarrow /a/$  in one case, and  $/@/ \rightarrow /i/$  in the other.

As introduced earlier, we first identify which vowel substitutions occur by auto-tagging a speech corpus based on acoustic characteristics. In order to determine the accuracy of the auto-tagger, we manually create a small labeled test set and evaluate the accuracy of the auto-tags against the manual labels. Once the auto-tagging process has been verified, we then tag a much larger corpus and determine the predictability of this auto-tagged corpus, using non-acoustic features.

In the remainder of this section, we discuss our approach to the auto-tagging process (Section III-A), the development of the manually labeled test set (Section III-B) and the feature selection and classification process (Section III-C). All experiments are performed using the SPCS corpus, introduced in section II.

#### A. Auto-tagging process

The SPCS corpus is first partitioned into a training and test set, and the training set is used to develop a standard Hidden Markov Model (HMM) based ASR system. The system is implemented using the HTK toolkit [12]. Acoustic models consist of cross-word tied-state triphones modelled using a 3state continuous density HMM. Each HMM state distribution is modelled by an 8-mixture multivariate Gaussian with a diagonal covariance matrix. The 39-dimensional feature vector consists of 13 static Mel-Frequency Cepstral Coefficients (MFCCs) with 13 delta and 13 acceleration coefficients appended. Cepstral Mean and Variance Normalization (CMVN) preprocessing is used and semi-tied transforms are applied.

During training, a pronunciation lexicon is used that retains the English schwa as a unit. Every word in this lexicon has a single variant: Sepedi words are modelled using Sepedi pronunciation models and English words using English pronunciation models. Sepedi pronunciations are obtained directly from G2P models (default-and-refine models [13] trained on the NCHLT *in-lang* dictionaries [14]). English

<sup>1</sup>The 'Speech Assessment Methods Phonetic Alphabet' is a standard computer-readable notation for phoneme descriptions. See [11].

words are first predicted using English G2P models[15], then manually reviewed. (As English words form the focus of the analysis, accurate pronunciations are expected to be important for the rest of the study.)

Once the ASR system has been trained, the same training data is realigned, using five different options: each replacing the schwa in the pronunciation lexicon with a different Sepedi vowel (/a/, /E/, /i/, /O/ or /u/). The resulting alignments then produce both timing information, as well as the likelihood of each vowel, given the data. These alignments are referred to as ' $spcs\_@$ ', ' $spcs\_a$ ', ' $spcs\_E$ ', etc. depending on which vowel was used during alignment.

Time alignments were manually verified to be accurate before proceeding. Accurate alignments were only obtained when training and aligning on the same corpus. Note that the lexicons are only changed during alignment, no re-training is performed. This ensures that all the data being studied are combined in the schwa-model, and not incorporated into any of the other vowel-models.

Once the likelihoods have been obtained, word-level GOP scores are extracted for each alignment option. Word-level GOP scores are used, as timing information otherwise introduces unnecessary variability. Word beginning and end times stay fairly consistent, while phone beginning and end times may differ significantly, especially if there is a mismatch between the vowel actually produced and the alignment candidate. For the same reason, frame-based (rather than phone-based) GOP scores are extracted, as defined in [10].

Two main results are obtained from the tagging process, for each schwa observed:

- The most likely vowel candidate (apart from schwa).
- Whether the phone matches the broad schwa category best, or is better matched to one of the Sepedi vowels.

These results form the basis of the analysis described in Section IV.

#### B. Manual tagging

In order to evaluate the effectiveness of the auto-tagging process, we created a manually labelled verification set. Two subjects were asked to listen to the words in question independently, and manually label each schwa with the most probable phone produced. Subjects were encouraged to only select 'schwa' if a true schwa was produced, and to otherwise select the closest vowel candidate. Both subjects were bilingual English/Afrikaans speakers, with subject B a trained linguist with exposure to Northern Sotho.

#### C. Classification process

Once the entire corpus has been labelled (using the autotagging procedure from section III-A), we review the results and identify possible features that influence vowel prediction. We use Naive Bayes classification to obtain an indication of whether these features are applicable. This is discussed further in Section IV.

#### IV. RESULTS AND ANALYSIS

We first analyse the validity of the manual tagging process (Section IV-A) before evaluating the accuracy of the autotagging process (Section IV-B). The influence of various possible non-acoustic features is evaluated in Section IV-D, before considering the overall predictability of vowel substitutions in Section IV-E.

#### A. Manual labels: inter-subject agreement

Any labels that were not true examples of code-switched schwas were removed from the evaluation set. (For example, where 'pressure' is only produced as / p r\E S / without articulating the last phoneme at all.) The number of remaining labels and inter-subject agreement measured using these labels, are shown in Table II. While inter-subject agreement is medium to fair at 70.0%, a review of the data shows that disagreement is mainly due to schwa boundaries being drawn in different places by the two subjects. If only labels that are not considered to be very close to schwa by either of the subjects are considered, inter-subject agreement increases to 85.7%.

TABLE II INTER-SUBJECT AGREEMENT DURING MANUAL TAGGING.

	#observations	#agreement	%agreement
All labels	50	35	70.00
Schwa excluded	35	30	85.71

#### B. Accuracy of the auto-tagger

As the auto-tagger is restricted to always select a possible alternative pronunciation (excluding schwa), it only makes sense to evaluate tagging accuracy against non-schwa labels. For the sake of completeness, results on both the full set and the non-schwa set are included in Table III.

#### TABLE III

ACCURACY OF THE AUTO-TAGGER, WHEN MEASURED AGAINST DIFFERENT MANUALLY LABELED TEST SETS.

	#observations	#agreement	%agreement
All labels: subject A	53	23	43.40
All labels: subject B	53	28	52.83
Schwa excluded: subject A	38	23	60.53
Schwa excluded: subject B	33	28	84.85
Schwa excluded: overall	71	51	71.83

From this analysis, it is clear that the auto-tagging process produces usable results, but agrees more strongly with subject B. While overall agreement is only at 71.8%, agreement between the auto-tagging process and subject B reaches 84.9%.

While the discrepancies may simply be due to greater transcription skill exhibited by subject B, it is worth understanding where human error occurs. For this reason, a formant figure was created as shown in Fig. 1 and Fig. 2.



Fig. 2. F1/F2 positions of labels. Each B/T legend displays the tag provided by subject B and the auto-tagger, respectively.

The first and second formants (F1 and F2) were extracted for each sample using Praat [16], and each sample plotted in the F1/F2 space. In Fig. 1, each sample is labeled with the tags from subject A and B. In Fig. 2, the sample is labeled with the tags of subject B and the auto-tagger. From the figures, it can be seen that most discrepancies among the three tags occur on the boundaries between classes. Tags in agreement are shown in green in both figures.

#### C. Corpus statistics

In total, 1 947 observations of schwa were auto-tagged, using the process described in Section III-A. The vowel /E/  $\!\!\!\!\!$ 

was tagged most frequently, and the vowel /u/ least. The number of tags associated with each vowel is shown in Fig. 3.



Fig. 3. Vowel distribution in the auto-tagged SPCS corpus.

#### D. Tag analysis

The tags were evaluated to determine how the observed vowel distribution is affected by non-acoustic factors. In Fig. 4 the vowel distribution is displayed per speaker. As is clear from this figure, speaker identity plays a minor role in determining which vowel is produced: per speaker, the distribution of vowels occurs approximately according to the same percentages as observed overall. This would immediately exclude other speaker-specific factors (such as gender or age). Speaker-specific features are therefore not considered further.



Fig. 4. The number of times each vowel was observed per speaker.

When analysing the tags per target word (all examples of the English word produced during code-switching are considered together), a clearer pattern emerges. The number of times a specific vowel is observed per word is shown in Fig. 5.

Based on the large role that word orthography plays, we next consider the graphemic string that produced the specific vowel (for example, -*a*-, -*i*-, -*io*- or -*ure*-). Note that the true prediction for each of these strings, in the specific context used, is 'schwa'. The data is first split into two parts to simplify analysis: where the schwa phone appears once in a word and where it appears multiple times. Since it is



Fig. 5. The number of times each vowel was observed per unique word.



Fig. 6. The number of times each vowel was observed per unique grapheme string.

expected that one schwa may be realised in different ways in the same word, the single-schwa words (where the schwa phone appears once in the corresponding phone string of the word) are analysed. In Fig. 6 we display the number of times each vowel is observed per unique grapheme string, for the single-schwa subset. We see that the graphemes *-ure-*, *-a-*, *-ia-*, *-ou-* and *-ur-* were mostly (but not exclusively) realised as phone */a/* during code switching. There are three graphemes *-e-*, *-io-*, *-er-* which were mostly realised as the Sepedi phone */E/*. There was very little confusion, for the graphemes *-o-* and *-i-* as they were almost always realised as phones */O/* and */i/*. The most unnpredictable grapheme string was -io-, which was realised as either */i/* or */E/*, two phones that overlap on the vowel chart in Fig. 1.

#### E. Predictability

From the above analysis, we select triphone and grapheme as input for a simple Naive Bayes (NB) classifier, in order to obtain an initial indication of predictability. Only the singleschwa words are considered. Using 10-fold cross-validation, we train models using only these non-acoustic features, and evaluate using the ten test partitions. Results are shown in Table IV. The rows show the phone label tags (obtained from the auto-tagger) and the columns are the label predictions from the NB classifier. We show the agreement level on the diagonal, both in terms of counts and percentage accuracy. An overall classification accuracy of 67.36% is achieved.

#### TABLE IV Confusion matrix when performing 10-fold cross-validation with non-acoustic features only.

	E	0	a	i	u
Е	214(73.5)	6	29	41	1
0	2	47 (68.1)	15	1	4
а	18	13	189 (84.8)	0	3
i	105	1	22	198 (60.7)	0
u	4	10	5	34	0 (0)

We see that the Sepedi phone labels E and i can be predicted fairly easily, although they are mutually highly confusable. On the other hand, the phone label u is never hypothesized. It also has the lowest occurance and predictability when analysed from a speaker-based, word-based, as well as grapheme-based perspective. It therefore seems better to not predict /u/ at all, but rather select either the secondmost probable candidate, or not to introduce a variant for words where /u/ is predicted as the most probable realisation. Which of these two strategies is better, will require further experimentation. As the grapheme string seems to be the main predictor of the matrix pronunciation of code-switched speech, it is expected that using the auto-tagging process to generate training data in order to train alternative G2P models, will be a productive area for further research.

#### V. CONCLUSION

Two interlinked processes were demonstrated: (1) a technique for auto-tagging an acoustic corpus, and (2) an analysis of the auto-tags to determine useful non-acoustic features for pronunciation prediction. These would make it possible to generate a lexicon for code-switched speech, prior to having encountered acoustic data related to the specific vocabulary to be modeled.

We found that the vowel-tagging process is quite difficult for humans to do, showing medium to fair inter-subject agreement if subjects are allowed to select the embedded phone itself (schwa in this case) as an option. When only those samples were included where speakers were deemed to have produced a substitute from the matrix language, inter-subject agreement increased to 85.7%. The auto-tagger agreed more strongly with one of the subjects: achieving a matching accuracy of 60.5% agreement with one, but 84.9% agreement with the other.

The non-acoustic features found to be most useful for predicting phone labels were triphone and grapheme string. Using these features, a simple Naive Bayes classifier achieves an average of 67.4% classification accuracy, evaluated using 10-fold cross-validation.

From this analysis we conclude that the best matrix pronunciations for the English schwa phone are E, i, O and a, with specific substitutions occurring based on the larger grapheme string. A one-to-one vowel substitution is only possible if the grapheme representation is either i or o. Using the trained classifier, pronunciation dictionaries can

now be developed by predicting the matrix pronunciations of English words in code-switched Sepedi speech, using only the vocabulary (no acoustic data) as input.

While the paper focused on the schwa found in English words embedded in Sepedi speech, the techniques are more generally applicable. (The English schwa is the foreign phone exhibiting the most variability, and was therefore selected as focus for this study.) Future work includes using a more sophisticated classifier, repeating this process for the full phone set, and conducting empirical ASR experiments using the proposed mappings.

#### VI. ACKNOWLEDGMENT

This work was partially supported by the National Research Foundation. Any opinion, findings and conclusions or recommendations expressed in this material are those of the author(s) and therefore the NRF do not accept any liability in regard thereto.

#### REFERENCES

- [1] T. I. Modipa, M. H. Davel, and F. de Wet, "Pronunciation modelling of foreign words for Sepedi ASR," in *Proc. Annual Symposium* of the Pattern Recognition Association of South Africa (PRASA), Johannesburg, South Africa, Dec. 2013, pp. 64–69.
- [2] T. Modipa, M. H. Davel, and F. de Wet, "Context-dependent modelling of English vowels in Sepedi code-switched speech," in *Proc. Annual Symposium of the Pattern Recognition Association of South Africa* (*PRASA*), Pretoria, South Africa, Nov. 2012, pp. 173–178.
- [3] C. Myers-scotton, Ed., Social motivations for Codeswitching: Evidence from Africa. Oxford: Clarendon Press, 1993.
- [4] Y. Li, Y. Yu, and P. Fung, "A Mandarin-English code-switching corpus," in *Proc. LREC12*, 2012, pp. 2515–2519.
- [5] D. Yu, L. Deng, P. Liu, J. Wu, Y. Gong, and A. Acero, "Crosslingual speech recognition under runtime resource constraints," in *Proc. ICASSP*, 2009, pp. 4193–4196.
- [6] V. B. Le, L. Besacier, and T. Schultz, "Acoustic-phonetic unit similarities for context dependent acoustic model portability," in *Proc. ICASSP*, 2006, pp. 1101–1104.
- [7] C.-L. Huang and C.-H. Wu, "Phone set generation based on acoustic and contextual analysis for multilingual speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP)*, 2007, pp. 1017–1020.
- [8] K. Sim and H. Li, "Context-sensitive probabilistic phone mapping model for cross-lingual speech recognition," in *Proc. Interspeech*, 2008, pp. 2715–2718.
- [9] S. Witt and S. Young, "Phone-level pronunciation scoring and assessment for interactive language learning," *Speech communication*, vol. 30, no. 2-3, pp. 95–108, 2000.
- [10] M. H. Davel, C. van Heerden, and E. Barnard, "Validating smartphonecollected speech corpora," in *Proc. SLTU*, Cape Town, South Africa, May 2012, pp. 68–75.
- [11] D. Gibbon, R. Moore, and R. Winski, *Handbook of standards and resources for spoken language systems*. Walter de Gruyter, 1997.
  [12] S. Young, G. Evermann, D. Kershaw, G. Moore, J. Odell, D. Ollason,
- [12] S. Young, G. Evermann, D. Kershaw, G. Moore, J. Odell, D. Ollason, V. Valtchev, and P. Woodland, "The HTK book," *Cambridge University Engineering Department*, vol. 3, p. 175, 2002.
- [13] M. Davel and E. Barnard, "Pronunciation prediction with Default&Refine," *Computer Speech and Language*, vol. 22, pp. 374–393, 2008.
- [14] E. Barnard, M. H. Davel, C. J. V. Heerden, F. D. Wet, and J. Badenhorst, "The NCHLT speech corpus of the South African languages," in *Proc. SLTU*, 2014, pp. 194–200.
- [15] L. Loots, M. Davel, E. Barnard, and T. Niesler, "Comparing manuallydeveloped and data-driven rules for P2P learning," in *Proc. Symposium* of the Pattern Recognition Association of South Africa (PRASA), Stellenbosch, South Africa, Nov. 2009, pp. 35–40.
- [16] P. Boersma and D. Weenink, "Praat: doing phonetics by computer (version 5.1.05) [computer program]," 2009, retrieved May 1, 2009, from http://www.praat.org/.

## Integration of an Electrical Discharge Machining Module onto a Reconfigurable Machine Tool\*

BH Roberts, IA Gorlach, TI van Niekerk Dept. of Mechatronics Nelson Mandela Metropolitan University Port Elizabeth, South Africa <u>Bryndan.Roberts@nmmu.ac.za</u>

Abstract—The objective of this research was to integrate an Electrical Discharge Machining module onto an existing reconfigurable CNC machine tool, using a modular approach, which potentially would enable conventional milling and EDM to be conducted in a co-ordinated fashion on the same machine tool. Ideally, multi-axis EDM is desired; however, single axis investigations were first done, as a step towards this ideal. The investigations determined which possible integration options should be pursued. This paper considers these single-axis investigations. The integration strategies implemented and evaluated in this research were (a) using the EDM machine's servo-drive output to drive the Reconfigurable Machine Tool's Z axis servo motor, via an external amplifier, (b) using a control signal from the EDM machine to communicate with the PCbased CNC control program (Mach3) to achieve advancement and retraction of the electrode using the RMT's Z axis servo motor, but governed by the EDM process, and (c) using an external controller, effectively replacing the EDM machine's controller and the CNC controller, while providing actuation of the electrode by means of a moving coil linear actuator mounted on the (now fixed) Z axis of the CNC machine. The external controller was implemented in LabVIEW, using a data acquisition device. The third strategy required that a voltage monitoring circuit be implemented to sense conditions at the EDM process spark gap. Single axis EDM was implemented using this method, but a version of the controller for potential 2dimensional machining was also developed. PC simulation models of the linear actuator, controller, and EDM process were produced.

Keywords—Electrical Discharge Machining; EDM; Reconfigurable Machine Tool; RMT; Reconfigurable Manufacturing Systems; RMS

#### I. INTRODUCTION

Over the past few years NMMU has developed a CNC Reconfigurable Manufacturing Tool (RMT) as part of a Reconfigurable Manufacturing System (RMS) [1,2]. RMS and RMT are defined in [3-6]. This machine tool is basically a 3-axis milling machine, but is intended to have a modular and reconfigurable nature. Fig. 1 shows the RMT. Control is achieved by means of a PC running Mach<sup>3TM</sup> [7] CNC software; a breakout board (interface device) provides step and direction pulses to the RMT's servo motors, via Gecko [8] servo motor drives.



Fig. 1. Reconfigurable 3-axis CNC machine tool [1]

The particular focus of the research involved the investigation of the integration of an EDM module onto the RMT, as a type of tool, such that the versatility and capability of the RMT was improved. Several such machines have been developed for research purposes [9], and example being the Generic Modular Machining Platform developed at Brunel University [10].

The EDM process, also known as spark erosion, is a particularly useful manufacturing technique often used in the tool, die and mould (TDM) environment, as very hard metals, e.g. hardened tool steels, can be machined easily. One reason for implementing EDM on a milling machine tool is that milling and EDM could be done in one set-up, i.e. without removing and resetting the workpiece. This would result in improved accuracy, and time saving. Also, stand-alone EDM machines (especially multi-axis machine) are very costly, whereas if one could add an EDM module (i.e. spark generator, with servo head if needed) onto a milling machine, costs could be reduced while functionality is maintained. Several such machines

#### II. EDM PROCESS AND EDM MACHINE AQUIRED

In the EDM process, high frequency voltage pulses from a spark generator are produced across a very small gap between an electrode and the workpiece; the resulting current pulses disintegrate the workpiece in the form of tiny pieces of debris, slowly eroding the workpiece surface [11-14]. The debris is removed by flushing of the di-electric fluid under which the

<sup>\*</sup>Sponsorship for this project from the Technology & Innovation Agency, and from the General Motors (SA) Chair of Mechatronics, is gratefully acknowledged.

EDM sparking takes place. A crucial aspect of an EDM process is the control system, whose main purpose it is to maintain the ideal gap distance, which it does by means of a servo driving system (conventionally a motor with leadscrew) for the electrode movement. A typical EDM feedback control system is shown in Fig. 2. Conditions at the spark gap are monitored and compared with the ideal conditions for the chosen machining parameters, which therefore effectively become a setpoint for the system. One method of monitoring the gap condition is by means of measuring the average voltage across the gap [13-15]. The gap voltage is strongly dependent on the gap distance; hence by measuring the voltage one is in effect measuring gap distance, which is one of the most important variables in the EDM process.



Fig. 2. EDM process feedback control system block diagram [13]

If the gap voltage is too high it implies that the gap is too large, and the servo-drive will advance the electrode; if too low, it will retract. Since the process has a noisy, random nature, the electrode is constantly advancing and retracting, in very small increments. It is clear that one cannot prescribe a feed-rate for EDM, as the gap voltage dictates the progression through the desired EDM path. I.e., the CNC controller provides the feed path, and the gap controller governs the instantaneous feed-rate and sense (i.e. advance/retract).

A suitable portable EDM machine, namely the Alic-1 produced by Sure-First of Taiwan, was sourced [16]. The Alic-1 comprises a servo head (electrode with servo motor actuator and lead screw), a spark generator (power pack), and a tank and pump for the dielectric fluid which must be present in the spark gap. The servo head can be mounted on a machine tool, usually by means of a magnetic base, for basic (manual) positioning of the workpiece relative to the EDM servo head. The power pack includes a front panel on which the operator can set and view certain parameters. Fig. 3 provides a basic control function diagram of the Alic-1 EDM machine's operation.



Fig. 3. Alic-1 EDM process control system block diagram

The encoder on the Alic-1 servo motor is not used during actual EDM; a varying voltage is simply applied to the motor armature to advance and retract the electrode, based on the voltage across the gap. It is however used for periodic 'jumping' of the electrode, to flush away debris from the gap.

#### III. INTEGRATION STRATEGIES

A simple integration of EDM onto a RMT would be to mount the EDM machine's servo head on the RMT's Z axis carriage; EDM would be by the EDM machine's servo head, but it could be combined with (independently controlled) RMT movements, mostly for XY plane positioning purposes, i.e. the RMT's axes would be used simply to move the EDM servohead into position above the workpiece, before commencing EDM. The servo head was in fact mounted on the Z axis, on a plate designed to attach to part of the spindle housing. Fig. 4 shows a picture of the servo head mounted on the RMT, while Fig. 5 shows the complete CNC machine combined with the EDM machine. Dowell pins were used for alignment of the servo head onto the plate, to ensure that the servo head could be easily removed and re-attached without losing the vertical alignment of the electrode, and without affecting the offset distances between the electrode and the milling spindle.



Fig. 4. RMT with EDM servo head attached



Fig. 5. Complete RMT incorporating EDM machine

Knowledge of the X and Y offset dimensions is important for co-ordination between the CNC machine and the EDM servo head, e.g. if one wants to use milling and EDM on the RMT without resetting the workpiece in-between the processes. A webcam was attached to the RMT under the spindle housing, viewing directly downwards (towards the workpiece), and this, together with a Mach3 'add-on' which allows webcam images to be displayed on the CNC PC screen, was used to establish the X and Y offset of the electrode relative to the spindle. Using these values, it is possible to specify tool offsets in the Mach3 tool setup screen, such that cross-over from milling to EDM could be achieved simply and accurately. The EDM machine could be made to communicate

with the CNC controller, for example to move (not while EDM is taking place) the servo-head laterally to new XY positions after a certain depth of EDM has been reached. This event would be known, since an encoder is present on the Alic-1's servo motor. Also, if for some reason very deep EDM was required, a signal could be sent to the CNC controller once the Alic-1's servo head reached the end of its travel - after which it could retract, the RMT's Z axis servo motor could advance, and then EDM could continue by means of the Alic-1's servo head. In this scheme the RMT's motors would not operate while EDM is taking place, and they would not be directly controlled by the EDM machine, i.e. they would be independent of spark gap conditions. The RMT's axis carriages would move only in-between EDM operations carried out by the servo head, and integration would be superficial.

Although integration in this sense would be a simple and practical solution for single-axis EDM, the main drawback is that it could not potentially accommodate multi-axis EDM. This is because a co-ordinated 3-axis feed path trajectory cannot be prescribed if one of the axes does not incorporate an externally accessible encoder, and a method of commanding the axis position at all times. Also, one may want to purchase a power pack (spark generator) separately, as a module, and still employ EDM on the RMT (e.g. by means of the RMT's motors). Hence this option was not pursued and is not considered as one of the investigated methods of integration. Testing of the electrode movement during EDM using the Alic-1's own servo head was however done, using a high accuracy displacement sensor, for later comparison with other integration techniques.

#### A. Using the EDM Machine's Servo-Drive Output to Drive the RMT's Z Axis Servo Motor

This control strategy was implemented mainly as a preliminary investigation of whether stable EDM could be achieved using the RMT's motors (with a view towards eventual multi-axis machining). A control block diagram for this implementation strategy is shown in Fig. 6, with the lower path having reference. The Alic-1's servo-drive output (a bipolar varying DC voltage), was amplified in order to drive the RMT's Z axis servo motor directly, using an independent amplifier (i.e. not the CNC Gecko servo-drive), while the electrode was fixed on the Z axis carriage (the servo motor in the Alic-1 servo head was disabled for this implementation method; i.e., the electrode was simply held in place on the carriage by means of the servo head housing). The RMT's Z axis servo motor is effectively controlled by the spark gap conditions (as sensed by the Alic-1 EDM machine).

Only single axis EDM could be achieved using this technique, as the relative speeds and positions of the axes would not be able to be controlled accurately, if the Alic-1's amplified servo-drive output was to drive two or more RMT axes simultaneously; but as mentioned, this implementation was seen as a preliminary step.

An H-bridge motor driver amplifier with PWM output was used for this integration method, with suitable signal conditioning between the Alic-1's servo-drive output and the amplifier. The EDM machine's servo-drive could not power the RMT's motor directly, due to its limited current and voltage ratings. In Fig. 6, either the EDM machine's servo head can be selected (i.e. for conventional operation), or the RMT's servo motor (Z axis) is used, together with the external amplifier and signal conditioning circuit.

The results of testing showed that, although EDM did occur, a large amount of hunting (over- and undershooting) occurred. The most likely reasons for the imperfect control are the high inertia (relative to the Alic-1's own servo motor and lead screw mechanism) of the Z axis carriage, and compliance due to the belt drive between the servo motor and the ballscrew. Also, the Alic-1's servo-drive controller parameters would not be tuned for the high inertia of the Z axis carriage. The other axes (lateral movement axes X and Y) involve even larger inertias, since the Z axis is carried on the Y axis, which is itself carried on the X (gantry) axis. Thus it would be futile to investigate EDM using the RMT's other axes using this methodology.

#### B. Using the Mach3 CNC Controller

In this strategy (depicted in Fig. 7 overleaf), which could



Fig. 6. Arrangement for use of Alic-1 servo-drive output, amplified to drive RMT servo motor



Fig. 7. Integration of EDM machine control and CNC RMT control

potentially accommodate multi-axis EDM, again the electrode was fixed on the Z axis carriage, i.e. the Alic-1's servo head motor was again disabled. The Z axis carriage was to be moved by the RMT's servo motors. However, in this scheme, Mach3 (CNC controller) would dictate the movement (step pulses) of the axes by virtue of G-Code commands, while the Alic-1's control system would modulate the execution of those movements (advance or retract) in accordance with the conditions at the spark gap. Multi-axis EDM could potentially be achieved using this integration method, because step pulse ratios for a desired linear travel could be determined in Mach3, according to the machining trajectory desired. Fig. 7 however depicts only a single axis implementation. Although this approach could suffer from the same high inertia problems already found, it was hoped that using this method may create more stable EDM, since the axis movements could effectively be slowed down (thus reducing overshoot) since the Mach3 controller provides timed step commands to the servo motor drives, and the instantaneous feed rate could be kept low.

Use was made of the polarity of the Alic-1's servo-drive output, as an input control signal for Mach3 (via the breakout board). The polarity of the Alic-1's output was sensed and this determined whether a G-Code command was given in Mach3 to advance, or to retract, the electrode (in small incremental steps). A so-called Macro written in VB Scripting language, which took into consideration the polarity of the signal, was used to pass G-Code commands to the main Mach3 program. The reason for using the Alic-1's servo-drive output as the input signal to Mach3 was that the gap condition signal internal to the EDM machine was not readily accessible.

Although a measure of control was achieved using this technique, the EDM process as a whole was not entirely satisfactory. It was found that although sparking did occur, it could not be maintained continuously, with frequent short circuits occurring.

#### C. Using External Controller and Moving Coil Linear Actuator

Although there were potentially advantages, in terms of stability, to using the CNC program with its step and direction output pulses (commanding the RMT's servo motor's position via its Gecko drive), control was found to be sluggish. The G-Code program suffers from feed rate ramp up and ramp down limitations, and minimum step increment size limitations. Also, when using a Macro to provide G-Code step commands, a delay (wait) function has to be employed to ensure that commands do not buffer before the G-Code program has time to execute them. In addition, a Gecko drive responds only to a position error (which takes a finite time to build up), whereas in certain control strategies velocity commands could be included, which would give a more immediate response.

Thus it was decided to investigate the use of an independent controller and servo motor driver, which could potentially give a more immediate response and be optimized for the task, since one would have access to all the control parameters. Also, since the RMT axes had been found to suffer from high inertia, it was decided to implement a low moving mass linear actuator instead of the RMT's servo motors.

Hence, the final implementation strategy made use of an external controller, implemented in National Instruments (NI) LabVIEW data-flow programming environment, and utilizing a data acquisition (DAQ) device between the PC running LabVIEW, an external H-bridge motor driver amplifier, and the encoder of a linear motor actuator (for electrode movement) acquired for the application. In this scheme, the RMT's CNC controller (Mach3) is completely bypassed, and movement commands would be generated by the external controller, if multi-axis EDM were desired (the controller could generate an electrode trajectory by employing multiple feedback loops simultaneously in a co-ordinated manner). A multi-axis EDM was physically implemented and is discussed here.

Although controller programs were developed for Proportional, Integral, and Derivative (PID) control, only one of a variety of control loops developed in this context is shown here, namely Proportional (P) control for a single EDM axis. Fig. 8 depicts a simplified block diagram of the control loop, while Fig. 9 shows a diagram of the physical implementation of the control loop. The input is the user-specified desired gap voltage (i.e. setpoint), and the feedback signal is the filtered



Fig. 8. Simplified block diagram of Proportional control system



Fig. 9. Block diagram of Proportional control system showing physical implementation

gap voltage. The difference (error) between these is the basis for the control action. The output of the program is a PWM and direction signal to the H-bridge powering the linear actuator motor armature. The PWM duty cycle is dependent on (proportional to) the gap voltage error.

The chosen actuator is a SMAC LCA25-010 moving coil device, which is basically a linear DC motor [17]. The actuator has a very low moving mass, and has a built-in encoder for shaft displacement sensing. The electrode was fixed directly to the actuator's moving shaft, while the housing was fixed to the Z carriage of the RMT, for manual pre-positioning purposes. If future multiple axis implementation was desired, additional linear actuators of the same type could be stacked such that the housing of one controller would be mounted of the shaft of another. Since the actuator housing is also of quite low mass, it is unlikely that high inertia problems would occur in multiple axis EDM. Compliance should also not be an issue since there are no lead screws or belts in the drive-train, and the lateral compliance of the actuator shaft in its housing is low.

The DAQ device used was a NI ELVIS (Educational Laboratory Virtual Instrumentation Suite) development board [18]. The PWM driver for the linear motor was based on the ST L298 dual full-bridge motor driver IC.

To achieve this method of integration, a gap voltage measurement and isolation circuit was implemented, rather than relying on the servo-drive output of the Alic-1 machine, since measurement of the gap condition provides a more direct signal. The average gap voltage was used in the control system, since the EDM process produces pulses. The pulses are also very noisy, due to the random nature of the sparking process, making filtering of the signal a necessity. The analogue signal was pre-filtered (not show in Fig. 9) using a toroid bead, and further filtering was implemented digitally in LabVIEW.

It is clear that with the particular control implementation depicted, namely armature voltage dependent on gap error, only single axis EDM is possible, since the actuator encoder does not form part of the control loop and the electrode position cannot be directly specified.

#### IV. MODELING OF CONTROL SYSTEM IN SCILAB

The Scilab model in Fig. 10 overleaf shows a basic position based control system, where a position error of the actuator (thus electrode) represents the EDM gap error, and random noise is applied to represent the EDM process. The Xcos environment of Scilab was used for modeling and simulation of the system, making use of Transfer Function (Laplace) representation. The main purpose was to model the actuator and controller, in order to potentially compare the model to the real implementation, and observe the effect of modifying various parameters, so that the process could in future be optimized.

In this model, the position setpoint is fixed, with the random noise superimposed. In reality, since the EDM process gradually removes material causing the surface height of the workpiece to reduce, the position setpoint itself effectively gradually changes, naturally. A version of the system where the erosion rate is summed into the model was developed, but is not shown here. The model shown does however provide a fair representation of the system, in terms of investigating control stability, since the erosion rate of material from the workpiece is orders of magnitude slower than the oscillatory movement of the electrode [15].

The innermost loop is part of the model of the actuator itself, which is taken from [19]. Some of the parameters for the Laplace representation of the actuator were found by experimental testing on the actuator, while others were estimated. Not all the parameters are visible in Fig. 10. In this model the position error controls the motor armature voltage by means of a PID function (primarily armature voltage is proportional to error, so this is similar to the Proportional control implementation discussed above).

A typical response graph for the model is shown in Fig. 11 overleaf, where it can be seen that although the electrode velocity (Green) fluctuates quite widely, the achieved position (Black) is relatively stable. The axis scales, and the graph values in general, are somewhat arbitrary because of the unknowns of the characteristics of the EDM process itself. The modeling was used in a qualitative rather than quantitative manner, where the main aim was to compare the effect on



Fig. 10. Scilab simulation model of actuator and EDM process



Fig. 11. Scilab model simulation output (step response)

stability of different types of controller. No attempt was made to quantitatively validate the model by direct comparison with the performance of the equipment, and this together with optimization of controller parameters is left to future work.

#### V. CONCLUSION

Varying results and levels of integration were achieved using the different approaches.

An H-Bridge motor driver was able to be used in conjunction with the Alic-1 EDM machine's servo drive voltage, to drive one of the machine tool's servo motors. EDM was achieved using this method, but control was poor, with a lot of electrode oscillation present.

The polarity from the Alic-1 EDM machine's servo-drive was able to be read into the machine tool's PC based Mach3 (CNC) program, and used to produce controlled advancement and retraction of the electrode, via the machine tool's Z axis servo motor. Smooth, continuous EDM could not be achieved using this method, although a certain amount of sparking did take place.

A third implementation technique utilized an external controller, together with a moving coil linear actuator. The controller was developed in LabVIEW using an NI ELVIS development board as a DAQ interface. This allowed filtering of the EDM gap voltage signal to be achieved, as well as the implementation of P and PID control. Stable single axis EDM was achieved.

The system architecture was able to be modelled in Scilab, in Laplace Transform block diagram form. The influence of the various controller types and parameters was able to be qualitatively observed in the response of the model.

#### REFERENCES

- I.A. Gorlach, B.H. Roberts, M. Simpson, W. Estment "Reconfigurable Manufacturing System for Mould and Die Making," PU-NMMU-NWU Consortium. 2009.
- [2] I.A. Gorlach, B.H. Roberts, M. Simpson, "Reconfigurable Manufacturing System for Mould and Die Making: Part Family Formation, Technologies and Process Planning Development," Progress Report for: PU-NMMU-NWU Consortium, 2010.
- [3] Artsoft, a division of Newfangled Solutions, [Internet], avaiable from <<u>http://www.machsupport.com/software/mach3></u> (Accessed 20 October 2015.)
- [4] Y. Koren, F. Jovane, T. Moriwaki, G. Pritschow, G. Ulsoy, H. Van Brussel H, "Reconfigurable manufacturing systems," CIRP Annals, 1999.
- [5] R. Katz, Y-M. Moon, "Virtual arch type reconfigurable machine tool design: principles and methodology," NSF ERC/RMS, University of Michigan, 2000.
- [6] M.G. Mehrabi, A.G. Ulsoy, Y. Koren, "Reconfigurable manufacturing systems: key to future manufacturing," Journal of Intelligent Manufacturing, 2000, 11(4):403-419.
- [7] R. Landers, B-K. Min, Y. Koren, "Reconfigurable machine tools," CIRP Annals, 2001, 49(1): 269-274.
- [8] CNC Direct, [Internet], Available from <<u>http://www.cncdirect.co.za/htm/servo.html</u>>(Accessed 6 Dec 2013.)
- [9] M.F. Devries, N.A. Duffie, J.P. Kruth, D.W. Dauw, B. Schumaker, "Integration of EDM within a CIM environment," CIRP Annals – Manufacturing Technology, 1990, 39(2):665-672
- [10] X. Sun, "An Integrated Framework for developing Generic Modular Reconfigurable Platforms for Micro Manufacturing and its Implementation," Brunel University School of Engineeirng and Design PhD Thesis, 2009, avaiable from <u>http://bura.brunel.ac.uk./handle/2438/3493</u> (Accessed 28 March 2013.)
- [11] M.P. Groover, "Fundamentals of Modern Manufacturing: Materials, Process, and Systems," Wiley, 2007.
- [12] E.J. Weller, editor, "Nontraditional Machining Processes," 2nd ed. Michigan: Society of Manufacturing Engineers, 1984.
- [13] K.H. Ho, S.T. Newman, "State of the Art Electrical Discharge Machining," International Journal of Machine Tools & Manufacture, 2003, 43:1287-1300.
- [14] M. Kunieda, B. Lauwrens, K.H. Rajurkar, B.M. Schumaker, "Advancing EDM through Fundamental Insight into the Process," CIRP Annals – Manufacturing Technology, 2005.
- [15] Y-G. Chang, "VSS Controller Design for Gap Control of EDM," JSME International Journal, 2002, Series C 45(3):712-721.
- [16] Sure First, [Internet], Available from <<u>http://www.surefirst.com.tw/html/02product\_07.htm</u>> (Accessed 6 December 2013.)
- [17] SMAC Moving Coil Actuators, [Internet], Product Flyer, LCA25 Series Linear Actuator, Available from <<u>http://www.smac-mca.com/pdf/LCA25-010%20flyer%20(web).pdf</u>> (Accessed Dec 2013)
- [18] National Instruments, NI.com, [Internet], NI ELVIS, Available from <<u>http://www.ni.com/ni-elvis/</u>> (Accessed 14 December 2013.)
- [19] B.C. Kuo, "Automatic Control Systems," Prentice Hall International Editions, 6th ed. 1991.

## **Text-based Language Identification of Multilingual Names**

Oluwapelumi Giwa<sup>1</sup> and Marelie H. Davel<sup>1,2</sup>

<sup>1</sup>Multilingual Speech Technologies, North-West University Vanderbijlpark, South Africa. <sup>2</sup>CAIR, CSIR Meraka, South Africa.

oluwapelumi.giwa@gmail.com, marelie.davel@gmail.com

Abstract-Text-based language identification (T-LID) of isolated words has been shown to be useful for various speech processing tasks, including pronunciation modelling and data categorisation. When the words to be categorised are proper names, the task becomes more difficult: not only do proper names often have idiosyncratic spellings, they are also often considered to be multilingual. We, therefore, investigate how an existing T-LID technique can be adapted to perform multilingual word classification. That is, given a proper name, which may be either mono- or multilingual, we aim to determine how accurately we can predict how many possible source languages the word has, and what they are. Using a Joint Sequence Modelbased approach to T-LID and the SADE corpus - a newly developed proper names corpus of South African names - we experiment with different approaches to multilingual T-LID. We compare posterior-based and likelihood-based methods and obtain promising results on a challenging task.

#### I. INTRODUCTION

Names form a large set of words with complex pronunciation patterns. For example, consider the name 'Wale', which can be a proper name in both Yoruba and English. This name will be pronounced completely different in Yoruba as / w { l e /, compared to /w e I l/ in English (using SAMPA<sup>1</sup> notation). Speech recognition and speech synthesis systems dealing with names require a means to produce accurate pronunciations for these words in order to function properly.

From a human perspective, different speakers tend to pronounce unfamiliar names taking factors such as orthographic irregularity and possible loanword status into account. Specifically, people tend to mimic pronunciations based on what they believe the origin of the name to be, or by replacing unfamiliar phoneme(s) with those they believe are the approximate in their mother tongue [2], [3].

In reality, there are numerous cases where documents or terms (either proper names or generic words) belong to more than one language of origin. Multilingual language identification (LID) is not a well-studied task, with existing textbased language identification (T-LID) techniques focussing either on identifying the single language within a portion of running text, or to a lesser extent, the language identity of isolated words. (See Section II.)

In practical terms, our question then is: given a proper name, can we accurately predict which source languages are relevant, where such names may be either mono- or multilingual? We specifically consider four South African languages using a newly developed proper names corpus. We build on previous work described in [4] using Joint Sequence Models (JSMs) for language tagging of isolated words.

The paper is structured as follows: Section II provides background on related T-LID studies. Section III describes the JSM-LID algorithm that we use as our basic technique. In Section IV, we provide an overview of our approach to the task of multilingual LID, extending the basic JSM-LID technique to cater for multilingual names. Section V presents the two different corpora used in subsequent experiments. Section VI describes our analysis and results. Finally, section VII summarises our findings.

#### II. BACKGROUND

Most research in T-LID focuses on the classification problem where a document belongs to one language category in a predefined language space. Documents can be either categorised as long- or short-text segment. According to [5], language identification of long text samples is a solved problem, in which approaches ranging from statistical to pattern recognition algorithms have been applied [6], [7], [8]. Recently, there has been renewed interest in language identification with the focus on short text segments [9], [10], [4], [11], [12], [13].

Text segments such as microblogging data, individual words in isolation and query search text are regarded as short. Traditional T-LID algorithms find this task a difficult one as less contextual information is available. Recent work in this area has especially been directed at microblogging domains [10], [14], [15], [11]. Bergsma et al. [11] examined LID on Twitter messages specifically for under-resourced languages, and found that systems trained on out-of-domain data obtained from Wikipedia outperformed other off-theshelf commercial and academic LID software (TextCat, GoogleCLD, Langid.py). They reported improved performance accuracy using compression-based language models of 97.0% (trained on Wikipedia), 97.4% using a maximum entropy classifier (trained solely on Twitter data), 97.9% using compression-based language models (trained on both Wikipedia and Twitter). Authors also made mention of the factors that contribute to higher performance accuracy such as training data, length of the tweet, and previous information across multiple tweets.

As the text becomes shorter, so the task becomes more difficult. Identification of isolated words (without context) has been approached using techniques such as dictionaries,

<sup>&</sup>lt;sup>1</sup>The 'Speech Assessment Methods Phonetic Alphabet' is a standard computer-readable notation for phoneme descriptions. See [1].

character *n*-gram language model, JSMs, support vector machines and conditional random fields [16], [4], [17], [18].

In previous work [4], JSM (Dependinga pronunciation modelling technique) was compared against SVM models for T-LID of words in isolation. Experiments conducted on four South African languages (Afrikaans, English, Sesotho and iziZulu) reported competitive results. The JSM-based system obtained an F1-measure of 97.2% as compared to a state-of-the-art SVM technique with an F1-measure of 95.2%. King and Abney [18] used a weakly supervised approach for identifying the languages of single words in a multilingual document. They experimented with different ranges of data sizes and reported that conditional random fields models trained with generalised expectation outperformed sequence classifiers.

From isolated words to proper name identification, task complexity increases. Konstantopoulos [19] examined language identification of proper names. He experimented with soccer player names obtained from 13 languages. He reported an initial average  $F_1$  score of 27% when tested on a general *n*-gram language model. With a more discriminated training data based on short sizes, an average  $F_1$  score of 50% was obtained on last names and 60% on first names. In related work, Li *et al* [20] used an *n*-gram language model to identify proper names in English, Chinese and Japanese. They reported an overall accuracy of 94.8% when classifying names amongst these three languages. (As these three languages are not closely related, the classification task becomes easier, explaining the high accuracy achieved.)

To the best of our knowledge, no studies are available that address the task of identifying multiple source languages of names in isolation. The closest related task addresses LID for multilingual documents (where a single document can belong to more than one language class). Approaches to this task include word-level identification [17], [18], vector-space model [21], minimum description length principle [22] and monolingual block segmentation [23], [24]. Within this context, Nguyen and Dogruoz [17] considered word-level classification in order to discriminate between Dutch and Turkish in a multilingual online discussion. They experimented with language model classification, logistic regression classification, dictionaries and conditional random fields; they reported the best performance using language models, while contextual information remains beneficial.

#### III. JOINT SEQUENCE MODELS FOR T-LID

In this section, we first describe the JSM algorithm and how it can be applied to T-LID, before extending it to the multilingual case in Sections IV and VI.

Joint Sequence Models were defined by Bisani and Ney in [25]. Initially developed for grapheme-to-phoneme (G2P) modelling, the technique is built around the concept of a 'graphone', an m-to-n alignment between small sections of graphemes and phonemes that form the basic units for probability modelling. Both the possible alignments and the graphones themselves are estimated through embedded maximization using a training dictionary. The probability of one unit occurring given the other(s) are similarly estimated using the same training data. The application of JSMs to LID is described in [4]. Below we provide an overview of the data preparation and training, as well as transcription phases. For more detail see [25] and [4].

#### A. Training phase

JSMs are typically used for pronunciation prediction. In order to generate a dictionary for T-LID training, all words with their corresponding language identifiers are added to a 'pseudo dictionary'. To simplify the co-segmentation task, a one-to-one letter to LID mapping is adopted whereby each grapheme corresponds to a single language identifier. For example, consider Table I where '#' represents word boundary markers, 'E' represents to English, 'Z' represents isiZulu, 'A' represents Afrikaans, and 'S' represents Sesotho.

TABLE I Example of a JSM dictionary recast for the T-LID task

Word	Pronunciation
#school#	EEEEEEE
#lekker#	ААААААА
#zuma#	ZZZZZZ
#lebo#	SSSSSS

As noted in [4], parameter choices such as graphone length, discounting, *m*-gram model order, and how models are initialised, influence performance accuracy of the JSM model. For this work, 1-1 graphones are trained, which means that the minimum and the maximum number of graphemes to LID mappings allowed per graphone length is 1. As models typically saturate before the  $8^th$  order, an  $8^th$ order model is used. Discounting is allowed during training to handle unseen token and avoid overfitting. All held-out data used for parameter estimation are folded back to the training set.

#### B. Transcription phase

In standard JSMs, a forward algorithm is used to compute the joint probability,  $p(g, \varphi)$ , of a co-segmentation between a grapheme sequence g and phoneme sequence  $\varphi$ . The probability of a source sequence g is in principle determined by summing all the matching graphone sequences across the sequence path. This value can also be approximated by simply taking the maximum value:

$$p(\varphi \mid g) = \frac{max_{q \in S(g,\varphi)}p(q)}{p(g)} \tag{1}$$

where  $S(g, \varphi)$  is the set of all co-segmentations of g, and p(q) represents the probability distribution over the sequences of graphones. (All probabilities as estimated during training).

JSMs allow the ability to generate different pronunciation variants. That is, for each word, a number of pronunciation variants may be produced; the path posterior probability given the word is used to select the winning candidate amongst the variants. As in the T-LID case, a variant consists of an LID string, ambiguity is resolved either by selecting the language identifier with the highest frequency counts
(within the variant) or by summing the language-specific log-probabilities internally to the word. (In Section VI-B, an alternative is proposed.)

#### IV. APPROACH

We approach the multilingual LID task by first selecting a technique that has been applied successfully to monolingual words in the past. Specifically, we select JSM-based LID, which demonstrates competitive performance for isolated word LID [4].

Two main issues must be addressed:

- 1) *The data to train the JSM models.* Is it better to use matching data (names) even if this data set is very small, or better to use a significantly larger set, even if it is unmatched (generic words)?
- 2) Options for extending the technique for multilingual *data*. JSMs can generate variant outcomes, as well as the likelihood and posterior probability given the word, per outcome. How well do these values predict true multilingual words?

We obtain empirical results by experimenting with different data sets and thresholds. We select the SADE multilingual name corpus [26] and the NCHLT *in-lang* dictionaries [27] for experimentation. These data sources are described in more detail in Section V. We discuss and experiment with different classification options in Section VI. We first propose an improvement to the monolingual classification technique, before experimenting with different approaches to multilingual classification.

As performance measures we use the standard definitions of precision, recall and F1-measure, that is:

$$precision = \frac{TP}{TP + FP}$$
(2)

$$\operatorname{recall} = \frac{TP}{TP + FN} \tag{3}$$

$$F1 = \frac{2 * \text{precision} * \text{recall}}{\text{precision} + \text{recall}}$$
(4)

with TP the number of true positives, TN the number of true negatives, FP the number of false positives and FN the number of false negatives. Note that 'recall' is the same as True Positive Rate (TPR). When a monolingual test set is used, these three measures are equivalent and we refer simply to accuracy.

When comparing threshold-based techniques we also evaluate the Receiver Operating Characteristic (ROC) by plotting the TPR (eq. 3) against the True Negative Rate (TNR) rather than against the False Positive Rate, as also often done, with

$$TNR = \frac{TN}{TN + FP} \tag{5}$$

In a ROC curve, the further the curve from the diagonal to the upper right hand side the better the performance of the system. Data points located far from the center of the origin depicts poor performance.

#### V. Data

We use two data sets in this analysis: the South African Directory Enquiries (SADE) corpus [26] to obtain tagged samples of multilingual names, and the NCHLT dictionaries [27] to obtain samples of generic monolingual words in matching languages. We analyse LID performance for four South African languages, namely Afrikaans, English, isiZulu and Sesotho.

The SADE corpus was developed to improve directory enquiry applications in South Africa. The corpus contains audio samples from multilingual speakers producing different proper names, and reflects a range of scenarios directory enquiries system could encounter. SADE data prompts were developed using publicly available names from a combination of Internet queries, personal names from a local tertiary institution (North-West University) and volunteers. Each name is annotated with a number of tags, including the most probable source language. We use only the word lists, and the LID tags of the v1.1. corpus in the current analysis.

The NCHLT dictionaries were developed in parallel with the NCHLT speech corpus [28]. For this work, we only use the word lists from the dictionaires, consisting of 15 000 unique words per language. These were estimated to be frequent words based on available corpus counts, and LID accuracy was verified using automated spell checkers and language practitioners [27]. We use the same edited lists as used in [4], where a second round of higher precision, lower recall spell-checking was performed for three of the languages (Afrikaans, isiZulu and English). These lists do not contain multilingual words.

#### VI. ANALYSIS AND RESULTS

We first analyse the distribution of multilingual words in the SADE corpus and create a suitable train and test set (Section VI-A) before applying JSMs to the monolingual test case (Section VI-B). The LID technique is then extended to cater for multilingual words in Section VI-C.

#### A. Data analysis

While the SADE corpus contains words in a large variety of languages, we only use the subset consisting of words tagged with the four target languages (Afrikaans, English, Sesotho and isiZulu). While the actual data set is almost exclusively bilingual by which Afrikaans - English definitely dominate. Per language, the number of unique words, average word length and total character count are displayed in Table II. The names contained in the corpus are not only person names, but also include the names of songs, restaurants and places, for example. These often consists of phrases (such as '*The Hillside Tavern*'). The resulting set of words is therefore a mixture of names and some generic words.

While the majority of the words are monolingual, a significant percentage of them (9.3%) were identified as multilingual. The distribution between mono- and multilingual words per language is shown in Table III. As before,

TABLE II	
SADE CORPUS: LANGUAGE DISTRIBUTION AND WO	ORD STATISTICS

Language	Word Count	Average word length	Character count
Afrikaans	1 050	6.9	7 308
English	6 634	7.4	48 733
Sesotho	465	7.8	3 612
isiZulu	458	8.1	3 689

the number of unique words, average word length and total character count are displayed.

#### TABLE III SADE CORPUS: MONO- AND MULTI-LINGUAL DISTRIBUTION AND WORD STATISTICS.

	Word	Word Count   Average word length   Character co			er count	
Language	Mono	Multi	Mono	Multi	Mono	Multi
Afrikaans	449	601	8.5	5.8	3 808	3 500
English	5 980	654	7.5	5.7	4 974	3 759
Sesotho	411	54	8.1	4.9	3 345	267
isiZulu	401	57	8.4	5.3	3 387	302

Due to the small word count of our isiZulu language, we randomly select a monolingual subset, per language, that equals to 401 unique words. For verification purpose, we ask language practitioners to review language tags of our newly extracted subset together with all the multilingual word-list shown in Table III. We analyse the results obtained from the language practitioners and observed few changes across words and their corresponding language IDs. In order to create a balance training set using only monolingual words, we randomly select a subset of 321 per language. Without restricting the allowed range of words per language in the test set, all remaining selection is random. The resulting training and test set partition is displayed in Table IV.

Most of the multilingual words in the corpus are bilingual, with a very small set of 3-lingual and a single 4-lingual word ('sale'). Examples of 3-lingual words include 'tutu', 'tone' and 'pole'. The exact distribution of words are shown in Fig. 1. Note that the figure is restricted to 800 words, even though the first bar exceeds this number (with a value of 7241 words).

The largest confusability exists between English and Afrikaans words, as can be seen from Table V, which lists the number of bilingual words in the SADE test set. For example, the word count of 314 represents the total number of words that do exist in Afrikaans and English only, 15 represents total word count that exist both in Sesotho and isiZulu only.

#### TABLE IV

NUMBER OF MONO- AND MULTILINGUAL WORDS IN THE SADE TRAIN AND TEST PARTITIONS.

	Traini	ng set	Test	set
Language	Mono	Multi	Mono	Multi
Afrikaans	321	-	155	329
English	321	-	137	348
Sesotho	321	-	98	34
isiZulu	321	-	99	34



Fig. 1. Number of 1-, 2-, 3- and 4-lingual words in the SADE full and test data sets.

TABLE V

LANGUAGE IDENTITIES OF BILINGUAL WORDS IN THE SADE TEST SET

Langu	Word count	
Afrikaans	English	314
Sesotho	isiZulu	15
English	isiZulu	13
English	Sesotho	7
Afrikaans	Sesotho	2

#### B. LID of monolingual names

For the generic model, we use a 40K-word subset of the NCHLT data. Our objective is to make use of an already existing model based on previous work done in [4]. For comparative purposes, the same statistics as shown for the SADE data in Tables II and III are displayed for the NCHLT data in Table VI. This provides us with two different training data sets: a small (1 284) SADE set, and a large (40K) NCHLT set.

TABLE VI NCHLT 40k subset: language distribution and word statistics.

Language	Word count	Average word length	Character count
Afrikaans	10K	10.5	104 531
English	10K	7.9	79 768
Sesotho	10K	8.3	82 537
isiZulu	10K	9.4	93 617

We first obtain results using the identical technique as described in [4]. Specifically, we sum the language-specific log-probabilities internally to the word. The latter technique (referred to as 'logprob sum') produced the most accurate results in [4]. We also report on an extension to this technique, where each variant is forced to be monolingual internally. That is, a 5-letter word will only be tagged as 'EEEEE' or 'ZZZZZ': a combination such as 'EEZZZ' is not allowed. In the 4-language task, this means each word would produce at most 4 variants, each with an associated posterior probability. These posteriors are then used to select the winning candidate, and any mixed variants simply discarded. This technique is referred to as 'forced pron' below.

#### TABLE VII

LID RESULTS FOR THE SADE MONOLINGUAL TEST SET, USING DIFFERENT TRAINING DATA SETS AND JSM-BASED TECHNIQUES.

Data set	Technique	Accuracy
NCHLT 40K	logprob sum	77.96
NCHLT 40K	forced pron	78.16
SADE	logprob sum	79.39
SADE	forced pron	81.02

In Table VII we report on results. Interestingly, the much smaller SADE training data clearly fits the test data better and produces more accurate results. The new 'forced pron' technique also shows a small but consistent improvement, across all measures.

#### C. LID of multilingual names

Using the SADE models and the 'forced pron' LID technique, we evaluate two approaches for determining when a word may be truly multilingual:

- We define an absolute threshold based on the posterior probability: any variants with a posterior probability higher than the threshold are accepted as additional source languages. We refer to this technique as 'absolute posterior' further.
- We define a relative threshold based on the log likelihood: any variants with a relative likelihood within the range of the best variant are accepted as additional source languages. We refer to this technique as 'relative likelihood' further.

In both cases, the best performing variant is automatically selected: thresholds are only used to determine whether more than one source language may potentially apply.

The ROC curves for both these methods are displayed in Fig. 2. (The closer the graph to the top right corner, the more accurate the technique.) The full test set from Table IV is used. As expected, the two techniques provide very similar results, with optimal  $F_1$  scores of 79.99% and 79.85% obtained, by 'absolute posterior' and 'relative likelihood', respectively.



Fig. 2. ROC curves for the SADE combined test set comparing the 'absolute posterior' and 'relative likelihood' approaches.

From the ROC curve in Fig. 2 it is clear that the techniques perform as expected, but how well do they perform? In order to answer this question, we consider a simple baseline whereby we select first the single best variant ('top-1'), and then the two best variants ('top-2'). In the first case, all words are therefore treated as if they are monolingual, and all additional source languages will automatically be 'false rejects'. In the second case, a large number of superfluous source languages will be hypothesized, and many 'false accepts' are accepted. We compare the results from these four methods in Table VIII.

TABLE VIII Comparing different multilingual classification approaches using the SADE combined test set.

Approach	Model	Recall	Precision	F-measure
top-1	NCHLT 40K	58.67	84.78	69.35
top-1	SADE	59.56	86.07	70.40
top-2	NCHLT 40K	91.89	66.51	77.17
top-2	SADE	91.98	66.84	77.42
relative likelihood	SADE	84.60	75.60	79.85
absolute posterior	SADE	84.80	75.70	79.99
relative likelihood	NCHLT 40K	86.00	77.60	81.58
absolute posterior	NCHLT 40K	86.20	77.80	81.79

From Table VIII it is clear that this is a difficult task: the first baseline result ('top-1') produces an F-measure of only 69.35% and 70.40% using the NCHLT 40K and SADE models, respectively. The second baseline approach ('top-2') shows improvement over the 'top-1' technique, where a recall value of 92% (obtained on both NCHLT 40K and SADE models) means a large number of the correct tags were correctly identified. However, this high recall value comes at a cost: specifically a precision that falls to 67%. As we accommodate more tags based on the number of variants, there is a high likelihood of also identifying the wrong labels. In contrast to this, the 'top-1' approach achieves better precision value, but with lower recall values.

We improve on the baseline results, where we leverage the trade-off between the 'false rejects' and 'false accepts' related to recall and precision respectively. For both approaches, we select an optimum point across different threshold values where TNR equals TPR, and compute the recall and precision values. The optimum threshold value where TNR equals TPR for 'relative likelihood' is -3.17, while the 'absolute posterior' is 0.026. Converting the optimum threshold value of 'relative likelihood' to a probability form, produces a value of 0.042, which is (as expected) close to the optimal value returned by the 'absolute posterior' approach.

Interestingly, we observe that using the NCHLT 40K model on 'absolute posterior' produces the best performance for multilingual classification of proper names. However, this contradicts our initial observation concerning monolingual classification, where the SADE model produced the best performance. This result shows that we cannot judge the performance of a multilingual classifier by only considering its monolingual classification accuracy.

#### VII. CONCLUSION

This work focussed on the multilingual classification of proper names, a task that has not been well-studied to date. Although our work is targeted at four South African languages, the techniques utilised are language independent.

First, we proposed an improvement to the monolingual classification of words using JSMs, building on earlier work described in [4]. By forcing JSMs to produce output strings that are associated with a single language, we obtain more trustworthy posteriors to analyse. We compared this approach to the best available to date ('logprob sum') and observed an improvement in F-measure from 79.39% to 81.02%, training and testing on the same set of proper names.

In order to classify proper names as multilingual, we experiment with two baseline methods ('top-1' and 'top-2') where both approaches produce a tradeoff between recall and precision. To strike a balance between our two metric values, we proposed two new techniques ('relative likelihood' and 'absolute posterior'). While the difference between the two new methods is statistically insignificant, both outperform the baseline with 'absolute posterior' using the NCHLT models producing an F-measure of 81.79%.

Finally, we observe that the identification performance on monolingual proper names does not necessarily translate to similar performance for multilingual classification. For the LID of monolingual names, models trained on SADE outperformed models trained on NCHLT 40K with an Fmeasure of 81.02% (compared to 78.16%). For LID of multilingual names, the NCHLT 40k models produced better results than SADE. (81.79% compared to 79.99%).

In conclusion, we have shown that even though LID of multilingual proper names is a challenging task, an adapted version of JSMs provide good classification accuracy.

#### VIII. ACKNOWLEDGMENT

This work was partially supported by the National Research Foundation. Any opinion, findings and conclusions or recommendations expressed in this material are those of the author(s) and therefore the NRF do not accept any liability in regard thereto.

#### REFERENCES

- [1] D. Gibbon, R. Moore, and R. Winski, *Handbook of standards and resources for spoken language systems*. Walter de Gruyter, 1997.
- [2] M. Benzeghiba, R. De Mori, O. Deroo, S. Dupont, T. Erbes, D. Jouvet, L. Fissore, P. Laface, A. Mertins, C. Ris *et al.*, "Automatic speech recognition and speech variability: A review," *Speech Communication*, vol. 49, no. 10–11, pp. 763–786, 2007.
- [3] F. Stouten and J.-P. Martens, "Dealing with cross-lingual aspects in spoken name recognition," in *Proc. IEEE Workshop on Automatic Speech Recognition & Understanding ASRU*, 2007, pp. 419–424.
- [4] O. Giwa and M. H. Davel, "Language identification of individual words with joint sequence models," in *Proc. 15th Annual Conference of the International Speech Communication Association, INTER-SPEECH*, 14-18 September, Singapore, 2014, pp. 1400–1404.
- [5] P. McNamee, "Language identification: a solved problem suitable for undergraduate instruction," *Journal of Computing Sciences in Colleges*, vol. 20, no. 3, pp. 94–101, 2005.
- [6] Y. Chen, J. You, M. Chu, Y. Zhao, and J. Wang, "Identifying language origin of person names with n-grams of different units," in *Proc. International Conference on Acoustics, Speech and Signal Processing ICASSP*, 2006, pp. 729–732.

- [7] C. Kruengkrai, P. Srichaivattana, V. Sornlertlamvanich, and H. Isahara, "Language identification based on string kernels," in *Proc. ISCIT*, 2005, pp. 926–929.
- [8] M. Padró and L. Padró, "Comparing methods for language identification," *Procesamiento del lenguaje natural*, vol. 33, pp. 155–162, 2004.
- [9] T. Vatanen, J. J. Väyrynen, and S. Virpioja, "Language identification of short text segments with N-gram models." in *Proc. of the 7th International Conference on Language Resources and Evaluation* (*LREC 2010*), Valetta, Malta, 2010, pp. 3423–3430.
- [10] E. Tromp and M. Pechenizkiy, "Graph-based n-gram language identification on short texts," in *Proc. of the 20th Machine Learning conference of Belgium and The Netherlands*, The Hague, Netherlands, 2011, pp. 27–34.
- [11] S. Bergsma, P. McNamee, M. Bagdouri, C. Fink, and T. Wilson, "Language identification for creating language-specific twitter collections," in *Proc. of the Second Workshop on Language in Social Media*, 2012, pp. 65–74.
- [12] M. Lui, J. H. Lau, and T. Baldwin, "Automatic detection and language identification of multilingual documents," *Transactions of the Association for Computational Linguistics*, vol. 2, pp. 27–40, 2014.
- [13] G. R. Botha and E. Barnard, "Factors that affect the accuracy of textbased language identification," *Computer Speech & Language*, vol. 26, no. 5, pp. 307–320, 2012.
- [14] S. Carter, W. Weerkamp, and M. Tsagkias, "Microblog language identification: Overcoming the limitations of short, unedited and idiomatic text," *Language Resources and Evaluation*, vol. 47, pp. 195–215, 2013.
- [15] M. Goldszmidt, M. Najork, and S. Paparizos, "Boot-strapping language identifiers for short colloquial postings," in *Proc. of the Machine Learning and Knowledge Discovery in Databases*, 2013, pp. 95–111.
- [16] O. Giwa and M. H. Davel, "N-gram based language identification of individual words," in *Proc. Annual Symposium of the Pattern Recognition Association of South Africa (PRASA)*, Johannesburg, South Africa, 2013, pp. 15–21.
- [17] D. Nguyen and A. S. Dogruoz, "Word level language identification in online multilingual communication," in *Proc. of the 2013 Conference* on Empirical Methods in Natural Language Processing (EMNLP 2013), Seattle, USA, 2014, pp. 857–862.
- [18] B. King and S. P. Abney, "Labeling the languages of words in mixedlanguage documents using weakly supervised methods." in *Proc. of the NAACL-HLT*, 2013, pp. 1110–1119.
- [19] K. Stasinos, "What's in a name? quite a lot," in Proc. of the 2007 Conference on Recent Advances in Natural Language Processing (RANLP-07), Borovets, Bulgaria, 2007.
- [20] H. Li, K. C. Sim, J.-S. Kuo, and M. Dong, "Semantic transliteration of personal names," in *Proc. of the Annual Conference of the Association for Computational Linguistics*, 2007, pp. 120–127.
- [21] J. M. Prager, "Linguini: Language identification for multilingual documents," in *Proc. of the 32nd Annual Hawaii International Conference* on Systems Sciences. HICSS-32, 1999, pp. 11–16.
- [22] H. Yamaguchi and K. Tanaka-Ishii, "Text segmentation by language using minimum description length," in *Proc. of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2012)*, Jeju Island, Korea, 2012, pp. 969–978.
- [23] W. J. Teahan, "Text classification and segmentation using minimum cross-entropy," in Proc. of the 6th International Conference Recherche d'Information Assistee par Ordinateur (RIAO' 00), 2000, pp. 943–961.
- [24] T. Mandl, M. Shramko, O. Tartakovski, and C. Womser-Hacker, "Language identification in multi-lingual web-documents," in *Proc.* of the 11th International Conference on Applications of Natural Language to Information Systems (NLDB 2006), Klagenfurt, Austria, 2006, pp. 153–163.
- [25] M. Bisani and H. Ney, "Joint-sequence models for grapheme-tophoneme conversion," *Speech Communication*, vol. 50, no. 5, pp. 434– 451, 2008.
- [26] Thirion, Jan W.F. and van Heerden, Charl and Giwa, Oluwapelumi and Davel, Marelie H., "The South African Directory Enquiries (SADE) corpus," in preparation.
- [27] M. H. Davel, W. D. Basson, C. van Heerden, and E. Barnard, "NCHLT Dictionaries: Project Report," Multilingual Speech Technologies, North-West University, Tech. Rep., May 2013. [Online]. Available: https://sites.google.com/site/nchltspeechcorpus/home
- [28] E. Barnard, M. H. Davel, C. J. V. Heerden, F. D. Wet, and J. Badenhorst, "The NCHLT speech corpus of the South African languages," in *Proc. SLTU*, 2014, pp. 194–200.

### A Local Planner for Ackermann-Driven Vehicles in ROS SBPL

Nicolas Limpert and Stefan Schiffer and Alexander Ferrein

Mobile Autonomous Systems & Cognitive Robotics Institute

FH Aachen University of Applied Sciences, Aachen, Germany

Email: nicolas.limpert@alumni.fh-aachen.de, s.schiffer@fh-aachen.de, ferrein@fh-aachen.de

Abstract—For a mobile service robot, safe navigation is essential. To do so, the robot needs to be equipped with collision avoidance and global path planning capabilities. The target platforms considered in this paper are Ackermann-driven vehicles. Hence, a planning approach which directly takes the kinematic and dynamic constraints of the vehicle into account is required. The Search-based Planning Library (SBPL) package of the Robot Operating System (ROS) provides global path planning which takes these constraints into account. However, it misses a local planner that can also make use of the Ackermann kinematic constraints for collision avoidance. A local planner is useful to take dynamic obstacles into account as early as possible. In this paper, we extend the SBPL included in ROS by a local planner, which makes use of motion primitives. We extend the ROS package and show first experimental results on Ackermann-driven vehicles.

#### I. INTRODUCTION

For a mobile service robot, safe navigation is essential. For this, the robot needs to be equipped with collision avoidance as well as with global path planning capabilities. The former is usually referred to as local path planning while the latter problem is called global path planning. The problem of path planning can be defined as in [1]:

Path planning is the problem of finding a path in the free configuration space  $CS_{free}$  between an initial configuration and a final configuration. If one could compute  $CS_{free}$  explicitly, then path planning would be a search for a path in this *n*-dimensional continuous space. However, the explicit definition of  $CS_{free}$  is a computationally difficult problem, theoretically (it is exponential in the dimension of CS) and practically. [...] Fortunately, very efficient probabilistic techniques have been designed that solve path planning problems even for highly complex robots and environments. They rely on the two following operations.

- 1) Collision checking checks whether a configuration  $q \in CS_{free}$  or whether a path between two configurations in CS is collision free, i.e., whether it lies entirely in  $CS_{free}$ .
- 2) Kinematics steering finds a path between two configurations q and q' in CS that meets the kinematic constraints, without taking into account obstacles.

Many service robots meet holonomic kinematic constraints (omni-directional drives), or can at least turn on the spot. Then, following a path in 2D space can, for instance, be done by decoupling the global path planning problem (finding a path between two points in the world) from the local path planning or navigation problem (collision avoidance). For instance, [2], [3] describe how to find a path in the geometric space of the robot, and then search for a realisation of that path in the configuration space of the robot. However, when it comes to car-like vehicles, one has to take the nonholonomic constraints into account. Then, decoupling the search for a path in Euclidean space from the search for a path in the configuration space cannot easily be done any more. Instead of planning a path in a discretized Euclidean space (grid), one can, for example, use lattice graphs based on motion primitives [4]. Motion primitives based on the possible configurations of the robot describe reachable states in Euclidean space. For a car-like drive as shown in Fig. 1(a), the parameters of the configuration space are the velocity and the steering angle. The robot moves on an arc which depends on both parameters. Different sets of parameters describe different arcs. To reach a goal in the global map, we need to find arcs consisting of different feasible parameters. Motion primitives encode different sets of these parameters and hence different arcs. The task is to find a sequence of feasible motion primitives that lead the robot to the goal.

In this paper we describe an extension of the SBPL planner, a lattice graph motion planning algorithm for carlike vehicles implemented in the Robot Operating System (ROS) [5]. The available implementation provides only a global planner, no local planner for Ackermann steering is available. This means, for one, that no local collision avoidance for car-like vehicle is available with SBPL. For another, if the robot fails to execute the pre-planned path due to slight imperfections (slip or drift), a new global plan has to be computed. In many cases, the global plan would still be valid, but a bridge plan to meet the global plan by making small adjustments would be required. Our contribution in this paper is:

- 1) a local planning algorithm based on motion primitives for car-like vehicles;
- an integration of the local planner into the ROS Move Base package together with the SBPL lattice planner used for global path planning.

We report on first experimental results of the SBPL global and local planner on a 1/10 model cart with Ackermann steering.

The remainder of this paper is organised as follows. In the next section, we review related approaches for robot motion planning. In Section III we introduce the SBPL planning algorithm and its implementation in ROS, before we describe



Fig. 1. The basic idea of lattice-based navigation planning based on motion primitives.

our extension to the planning system in Section IV. It is followed by first experimental results in Section V. We conclude with a discussion and an outlook to future work.

#### II. RELATED WORK

As described in the introduction, there are approaches that decouple the path planning and the motion planning problem. One such example is Stachniss and Burgard [2]. In their approach, first a Euclidean path is found in the global map. To follow this path with the robot, valid motion commands in an area around the Euclidean path are sought. Our approach follows a similar idea. It first renders a feasible path in the Euclidean space around the robot. It then tries to follow this path as close as possible. Compared to their planning approach, ours rather follows the idea of Pivotraiko and Kelly [6] which made use of pre-computed motion primitives in order to plan with feasible motions and to speed up planning. Their work is not restricted to Ackermanndriven robots but rather provides a general method to plan smooth trajectories. Our approach can also be used for other robotic platforms-one only has to provide the several motion capabilities in the form of motion primitives and tune the appropriate executing component (e.g. one has also to consider sideways movement for holonomic robots).

The idea of planning motions with motion primitives can also be transferred to Bezier curves which let the user plan smooth motions for non-holonomic robots. In [7], Liang et al. present an approach to plan motions for an automatic parking assistant based on Bezier curves and a PID controller. The approach is able to plan smooth trajectories for the parallel parking tasks for Ackermann-driven robots. The first instance of our planner makes a similar use of a PID approach in order to get the current difference from the robot's position to the next way-point on the path. Rather than using motion primitives, Bezier curves can probably serve motion plans feasible for the particular robot in narrow situations. They can lead to better motions than the combination of motion primitives. The downside is that every motion combination has to be calculated from scratch rather than making use of pre-computed motion primitives.

Choi et al. [8] also had the idea of making use of piecewise Bezier curves to combine a set of these curves as a set of feasible motions in an appropriate order. The usage of piece-wise Bezier-curves is much closer to the idea of a set of motion primitives which are used by SBPL. The Bezier curves are parameterised according to the motion constraints of the robot. It is therefore an interesting work for the SBPL Lattice Planner as it takes the idea of motion primitives further by defining the different dynamic constraints not as motion primitives, but mathematically as Bezier curves.

clarity, no nodes at 45° are expanded.

An even more complex motion planning problem is tackled in the work of Laumond [9]. This work incorporates motion plans for articulated vehicles. The approach is to generate vector fields in order to achieve feasible motions. This idea is more complex than the generation of motion primitives as the planning approach has to take other body shapes into account. However, it makes an interesting alternative as it lets the robot plan in much more dynamic environments just like the Bezier curves do. Our approach is founded on the work of Likhachev et al. [10]. They make use of lattice graphs based on motion primitives for selfdriving cars. They implemented their approach in the ROS SBPL Lattice Planner which is presented in detail in the next section. Our work is founded on this approach.

Making use of motion primitives is very similar to the work of F. von Hundelshausen et al. in [11]. Their approach is to incorporate a 360 degree laser range scanner in order to evaluate a set of "tentacles". These "tentacles" form a set of feasible motions based on the current robotic vehicle's velocity and surrounding obstacles. This makes them a very similar idea to the motion primitives used in our approach. Their benefit is that they can incorporate the current robot's velocity. Another work that takes Ackermann geometries into account is [12] which makes use of Dubin's paths [13]. The idea is to perform motion commands that only consist of a combination of the plain motions left, forward and right. Reeds and Shepp [14] extend this work and add backward movements. The benefit of this approach over ours is that it always delivers the optimal path under the assumption that there are no obstacles in the way. Unfortunately how-



Fig. 2. SBPL Lattice Planner planning through a corridor with too sharp a turn for the local planner.

ever, in a real world scenario there are obstacles in the way. In the ROS environment there exists a local planner called ackermann\_local\_planner which makes use of Dubin's Paths. This work has not been considered in our experimental setup for a comparison as it does not provide any collision avoidance at all.

In ROS, a navigation package called the move\_base which makes use of a global and a local planning aspect is provided. The move\_base also delivers a so-called base\_local\_planner which provides implementations of the Trajectory Rollout Planner and Dynamic Window Approach. They both deliver the idea of sampling a set of motions and choose the appropriate motion in order to get a plan to reach a navigational goal. However, the base\_local\_planner is used for robots that are at least differentially driven and can rotate in place. This makes it difficult to be used on Ackermann-driven vehicles as they cannot rotate in place.

#### III. THE SBPL PACKAGE FOR ROS

In [10], Likhachev and Ferguson describe a global planner for generating complex dynamically-feasible manoeuvres for autonomous vehicles travelling at high speeds over large distances. They use an approach by Pivotraiko and Kelly [15] to generate lattice graphs based on motion primitives. Their approach makes use of the anytime A\* variants ARA\* and AD\* [16] to come up with a dynamically feasible motion plan at any time given. When travelling at high speeds, it is important to have an immediate answer what (dynamically feasible) action the robot should perform next in order to reach the next target point. The SBPL approach was used by the Tartan Racing Team [17], the winning team of the DARPA Urban Challenge (DUC). For a concise overview of the DUC, see [18].

The basic principle of lattice graphs based on motion primitives is shown in Fig. 1. Fig. 1(b) shows motion primitives for a robot with 5 different successor states. Based on the kinematic and dynamic constraints of the robot platform and the current state, robot platforms with non-holonomic motion constraints only have certain feasible successor states. For instance, as a non-holonomic robot cannot turn on the spot, it has to perform quite some manoeuvre (driving forwards, pushing backwards, driving forwards again) to come to a stand-still at the same position but with a different orientation depending on the turning angle of the vehicle. The parameters for the Ackermann drive (Fig. 1(a)) are steering angle and translational speed. Based on these two parameters and the current state, a feasible successor state of the vehicle which is reachable from the current state can be computed.

Now, when trying to navigate to a point in the global frame of the robot, one has to derive the feasible drive commands to reach that very point. Unlike road-map approaches which first compute a path in the 3D space and then generate the motion commands, lattice graph planning approaches already take the feasible commands into account while planning a path to the goal. Fig. 1(c) shows an example for our 5-motion primitive. In each state, feasible motion primitives can be applied to come up with the successor state in the search graph. The nodes of the graph represent the pose of the robot in 3D space, the arcs denote feasible motion commands.

The SBPL Lattice Planner is an implementation of a global planner using motion primitives in ROS.<sup>1</sup> It can be used as a global planner within the move\_base package.<sup>2</sup> The ROS move\_base package is a navigation package for global and local path planning algorithms. It takes standard information such as odometry information, a global map or a localisation pose into account to come up with a cmd\_vel message that is directly sent to the motors. The cmd\_vel message gives the translational and rotational velocity the robot has to execute in the next time step. In addition, this package implements a state machine for recovery behaviours and a local cost map. The cost map is a local map around the robot which takes dynamical obstacles (which are not represented in the global map) into account. The local planner uses this cost map for collision avoidance while following way points computed on the global map. The cost map can also incorporate different dynamical layers which can include environment information gathered from the robot's sensors.

Figure 2 shows an example of a global plan derived by SBPL for a narrow passage, where SBPL tries to find a solution through this passage. The robot is located at the left end of the corridor, the target point is around the opposite corner. The black cells represent the global cost map, i.e. the obstacles that could be perceived by the sensors of the robot. The small white dots represent the sensor readings from a laser range finder showing the walls of the corridor. One can observe that the obstacles and walls are extended in the cost map to guarantee a collision-free path. One can further observe the application of the motion primitives leading to a left-right-left combination of turns to navigate around the corner at the end of the corridor. The final left, however, touches an obstacle in the global cost map. The SBPL Package for ROS is also used for other planning tasks, for instance, for the manipulator planning package MoveIt<sup>1</sup>.<sup>3</sup>

#### IV. A LOCAL SBPL PLANNER FOR CAR-LIKE VEHICLES

In order to plan feasible motions for our Ackermanndriven robots (Fig. 3) we deployed the ROS SBPL planning package as a local planner and merged it with the ROS pose\_follower package for local navigation. Addition-

```
<sup>1</sup>http://wiki.ros.org/sbpl_lattice_planner
<sup>2</sup>http://wiki.ros.org/move_base
```

```
<sup>3</sup>http://moveit.ros.org/
```



(a) FHAC Rover is a 1/10 car model.

(b) Rugged-terrain vehicle MQOne with Ackermann steering [19].

Fig. 3. Ackermann-driven robots used in the experiments.



(b) Re-planning in local frame

Fig. 4. Comparison of global and local re-planning on sudden obstacle

ally, we use the carrot\_planner<sup>4</sup> to get very basic global goals (Euclidean description of start and goal).

As we mentioned in the previous section, the SBPL planner works fine for global path planning. It takes the dynamic constraints of the robot platform into account to come up with a plan consisting of feasible motion patterns of the used vehicle. One drawback, however, is that the SBPL is only available as a global planner. This means that dynamic obstacles that appear right in front of the robot and hence are not represented in the global map of the robot, are not taken into account. For the original use-case of the package, this was not a problem.

Our use-cases, however, are a bit different, since our

vehicles drive at much lower speeds and in more narrow environments. We want to use the SBPL planner for our 1/10 FH Aachen Rover (Fig. 3(a)) as well as for our rugged-terrain mining vehicle MQOne (Fig. 3(b)). The former vehicle is a 1/10 model of a car-like vehicle which is used for teaching activities such as the International ROS Summer School [20]. The second vehicle is the research robot MQOne [19] which is a rugged-terrain robot design used for mapping tasks in underground mines. In our target scenarios we need to take dynamic obstacles into account and we should take such obstacles into account as early as possible. This is especially useful for Ackermann-driven platforms to avoid unnecessary manoeuvring.

In our approach, we use SBPL for local planning. Once the global planner (carrot\_planner) sends the goal, the local planner receives a list of way-points. The local planner then considers the current deviation from the robot's position to the next global way-point and plans a path to that respective way-point by making use of the SBPL Lattice Planner. As shown in Fig. 1, SBPL combines the possible motion primitives in order to reach the desired goal position as closely as possible. When the local planner generated a path to the next way-point, the pose\_follower (basically a P controller for linear and angular velocities) generates motion commands to drive to the particular way-point. If an obstacle is detected along the path, local re-planning is initiated. This way, we consider dynamic obstacles as soon as possible and only re-plan to the next global way-point instead of re-planning globally. In Fig. 4(a), the result of a global re-planning of the path is shown.

- 1) The robot realises that it is about to collide with a dynamic obstacle.
- 2) It starts some emergency stop behaviour.
- 3) It tries to find a global path around the obstacle.

In Fig. 4(b) our novel SBPL-based local planner can be seen in action.

1) The ROS cost map is updated with the dynamic

<sup>&</sup>lt;sup>4</sup>http://wiki.ros.org/carrot\_planner

obstacle.

- 2) The local planner initiates the planning procedure to find motion commands that guide the robot around the obstacle back on the globally planned path.
- 3) The advantage is that the global plan is still valid after surrounding the obstacle.

Hence, no new global planning needs to be initiated when collision avoidance kicks in. Further, using motion primitives for trying to find a way around the obstacle leads to smoother overall behaviour of the robot. One should note that an emergency stop needs to be in the portfolio for the local planner as well to ensure that the robot can stop in front of the obstacle in case no local plan around the obstacle could be found. The test scenario shown in Fig. 4 was simulated using the ROS simulation environment Gazebo.

#### V. FIRST EMPIRICAL RESULTS

In this section we show first experimental results. We compare our novel local SBPL-based planner with SBPL as a global planner in combination with the aforementioned pose\_follower for local navigation. We compare both approaches in terms of planning time, path length, and number of re-planning steps in different simulation scenarios. For simulating the 1/10 model car, we use the Gazebo simulator.<sup>5</sup> The tasks in each of the scenarios were to advance about 10 meters around different obstacles. These obstacles were unknown in the initial planning phase. The first test was to try and avoid a pylon. In the second test scenario, the robot drove directly towards a wall standing in the environment. In the third test scenario, the robot had to drive around a slalom course where several blocks were obstructing the direct path to the goal.

Results are shown in Tab. I. The numbers result from averaging over twenty runs for every scenario. The planning time refers to the time that the SBPL component in each configuration needs for planning. In the scenario, where the robot was to drive around a pylon, both approaches are nearly equal in terms of the total travel distance and planning time. The local planner has a slightly lower path execution time despite the fact that it re-planned two times more often than the global planner.

In the second scenario (Tab. I(b)), where the robot was heading towards a wall and had to plan a path around it, our novel local SBPL planner was able to execute the planned path much faster. The pose\_follower needed 50 seconds more despite a similar path length and planning time. The advantages of using a local planner in combination with a global planner over using just a global planner can clearly be seen in our third test scenario. The results are shown in Tab. I(c). This was a test setup with several blocks standing in front of the rover forming a slalom course. The robot had to plan a way through the blocks to reach the target point. Our approach yielded a smaller travel distance and a much shorter execution time. The execution time of the our approach is only half of the global planner plus the pose follower. The local planner re-planned about 4 times more often than the global approach, yielding nonetheless a shorter overall planning time.

These results suggest that using a local SBPL approach together with a global one is beneficial in general. One of the reasons is that with the global planner plus pose follower re-planning is only initiated when the robot decides that the global path cannot be followed as it is leading through an obstacle.

Initial testing was in simulation, secondary testing was with real hardware. Fig. 5 shows a city-like environment that we used during our International ROS Summer School which we held between August 10–24 in Aachen. The challenge for the students from over 20 nations was to program a pizza delivery robot. They made use of a number of different ROS packages such as Hector SLAM for mapping and localisation [21]. Fig. 5(c) shows the cost map of the environment consisting of four blocks surrounding drive-ways. Fig. 5(d) shows a planned path in this environment. While some more tweaking of the motion primitives would be required to increase the planner's performance, we demonstrated successful navigation of Ackermann-driven vehicles through narrow passages.

#### VI. DISCUSSION

In this paper we presented a method for navigation on Ackermann-driven mobile robots where we use SBPL as a local planner. This alleviates incorporating the kinematic constraints while handling dynamic obstacles. SBPL uses motion primitives and hence considers the kinematic constraints. This is especially useful for Ackermann-driven vehicles as the constraints might result in quite extensive manoeuvres. Using SBPL as a local planner improves handling of dynamic obstacles as compared to only using SBPL as a global planner. This is because the global planner considers nonstatic obstacles only when collision is immanent, whereas the local planner can modify the path as soon as an obstacle comes into the field of view. In our experimental evaluation we show that using our local planner is successful in both, a Gazebo simulation environment and in a real scenario with the FHAC Rover.

Planning with motion primitives has the benefit of being able to define the particular kinematic constraints in order to plan feasible motions to reach a navigational goal. Providing a component for local navigation that uses motion primitives, is improving the handling of dynamic obstacles. This is of particular importance if the kinematic constraints cause cumbersome manoeuvring in order to react to local changes like a nearby obstacle. In our evaluation we could show that in such situation we can cut the time needed to complete a path in half, using SBPL as a local planner.

The downside of motion primitives is that one always has to evaluate and define the capabilities of the particular robot in the form of motion primitives manually. If one could let the robot find out its very own physical constraints and define those as motion primitives it would enable the robot to perform planning based on the particular motion constraints.

<sup>5</sup>http://gazebosim.org/

#### TABLE I Results from different test scenarios

(a) T	he pylon scenari	0.	(b) The wall scenario. (c) The slalom course sce			nario.		
Data	global SBPL +	carrot +	Data	global SBPL +	carrot +	Data	global SBPL +	carrot +
	pose_follower	local SBPL		pose_follower	local SBPL		pose_follower	local SBPL
travel dist [m]	11.342	9.342	travel dist [m]	16.837	10.698	travel dist [m]	18.986	12.945
execution [s]	40.839	35.493	execution [s]	90.8	41.6	execution [s]	90.4	48.5
planning [s]	0.196	0.189	planning [s]	0.365	0.331	planning [s]	0.343	0.314
avg. # re-plans	2.9	5.65	avg. # re-plans	3.9	4.6	avg. # re-plans	6	17.7



(a) City-like environment (streets and blocks) used in the ROS Summer School 2015.

(c) Map of the environ-(d) Screenshot of path in exement cution

Fig. 5. Planner in action at the ROS Summer School 2015

#### ACKNOWLEDGEMENTS

We thank the ROS Summer School team for their support in this work. We would further like to thank the anonymous reviewers for their valuable comments and suggestions to improve the quality of the paper.

#### REFERENCES

- M. Ghallab, D. Nau, and P. Traverso, Automated Planning: Theory & Practice. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2004.
- [2] C. Stachniss and W. Burgard, "An integrated approach to goaldirected obstacle avoidance under dynamic constraints for dynamic environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland, September 30 - October* 4, 2002, 2002, pp. 508–513.
- [3] S. Jacobs, A. Ferrein, S. Schiffer, D. Beck, and G. Lakemeyer, "Robust collision avoidance in unknown domestic environments," in *RoboCup* 2009: Robot Soccer World Cup XIII, ser. LNCS, vol. 5949. Springer, 2009, pp. 116–127.
- [4] M. Pivtoraiko and A. Kelly, "Efficient constrained path planning via search in state lattices," in *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, 2005.
- [5] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009.
- [6] M. Pivtoraiko and A. Kelly, "Kinodynamic motion planning with state lattice motion primitives," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2011, pp. 2172–2179.
- [7] Z. Liang, G. Zheng, and J. Li, "Automatic parking path optimization based on bezier curve fitting," in *IEEE International Conference on Automation and Logistics (ICAL)*. IEEE, 2012, pp. 583–587.
- [8] J.-W. Choi, R. Curry, and G. Elkaim, "Piecewise bezier curves path planning with continuous curvature constraint for autonomous driving," in *Machine Learning and Systems Engineering*. Springer, 2010, pp. 31–45.
- [9] J.-P. Laumond, "Nonholonomic motion planning versus controllability via the multibody car system example," DTIC Document, Tech. Rep., 1990.
- [10] M. Likhachev and D. Ferguson, "Planning long dynamically feasible maneuvers for autonomous vehicles," *The International Journal of Robotics Research*, vol. 28, no. 8, pp. 933–945, 2009.

- [11] F. Von Hundelshausen, M. Himmelsbach, F. Hecker, A. Mueller, and H.-J. Wuensche, "Driving with tentacles: Integral structures for sensing and motion," *Journal of Field Robotics*, vol. 25, no. 9, pp. 640–673, 2008.
- [12] X.-N. Bui, J.-D. Boissonnat, P. Soueres, and J.-P. Laumond, "Shortest path synthesis for dubins non-holonomic robot," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1994, pp. 2–7.
  [13] L. E. Dubins, "On curves of minimal length with a constraint on
- [13] L. E. Dubins, "On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents," *American Journal of mathematics*, pp. 497–516, 1957.
- [14] J. Reeds and L. Shepp, "Optimal paths for a car that goes both forwards and backwards," *Pacific journal of mathematics*, vol. 145, no. 2, pp. 367–393, 1990.
- [15] M. Pivtoraiko and A. Kelly, "Generating near minimal spanning control sets for constrained motion planning in discrete state spaces," in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (*IROS*). IEEE, 2005, pp. 3231–3237.
- [16] M. Likhachev, D. Ferguson, G. Gordon, A. Stentz, and S. Thrun, "Anytime search in dynamic graphs," *Artificial Intelligence*, vol. 172, no. 14, pp. 1613–1643, 2008.
- [17] D. Ferguson, T. M. Howard, and M. Likhachev, "Motion planning in urban environments," *Journal of Field Robotics*, vol. 25, no. 11-12, pp. 939–960, 2008.
- [18] M. Buehler, K. Iagnemma, and S. Singh, *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*, 1st ed. Springer Publishing Company, Incorporated, 2009.
- [19] S. Rebel, F. Hüning, I. Scholl, and A. Ferrein, "Mqone: Low-cost design for a rugged-terrain robot platform," in *Intelligent Robotics and Applications - 8th International Conference, ICIRA 2015, Portsmouth, UK, August 24-27, 2015, Proceedings, Part II, ser. LNCS, vol. 9245,* 2015, pp. 209–221.
- [20] A. Ferrein, S. Kallweit, I. Scholl, and W. Reichert, "Learning to program mobile robots in the ros summer school series," 2015.
- [21] S. Kohlbrecher, J. Meyer, T. Graber, K. Petersen, U. Klingauf, and O. von Stryk, "Hector open source modules for autonomous mapping and navigation with rescue robots," in *RoboCup 2013: Robot World Cup XVII*, ser. LNCS. Springer Berlin Heidelberg, 2014, vol. 8371, pp. 624–631.

# HMM Adaptation for child speech synthesis using ASR data

Avashna Govender<sup>1</sup>, Bakari Nouhou<sup>2</sup> and Febe de Wet<sup>1</sup>

Abstract-Acquiring large amounts of child speech data is a particularly difficult task. One could therefore consider the possibility to add existing corpora of child speech data to the severely limited resources that are available for developing child voices. This paper reports on a feasibility study that was conducted to determine whether it is possible to synthesize good quality child voices using child speech data that was recorded for automatic speech recognition (ASR) purposes. A text-tospeech system was implemented using hidden Markov model based synthesis since it has proven to be a technique that is less susceptible to imperfect data. The paper describes how data was selected from the ASR corpus to build various child voices. The voices were evaluated to determine whether the data selection methods yield acceptable results within the context of model adaptation for child speech synthesis. The results show that, if data is selected according to particular criteria, ASR data could be used to develop child voices that are comparable to voices that were built using speech data specifically recorded for speech synthesis purposes.

#### I. INTRODUCTION

Child speech synthesis is considered to be a daunting task due to the difficulty of collecting appropriate data. Numerous challenges are associated with collecting speech data from children. For instance, children usually have a very short attention span and as a result recording sessions need to be broken down into shorter sessions over longer periods of time [1]. Younger children's speech is characterized by many inconsistencies such as stuttering and repetition because they are beginner readers. Children also tend to become emotionally involved in the text they read which leads to fluctuations in their speech associated with emotional expression. These are all characteristics that are not favourable for the type of speech needed to train a synthetic voice [2].

For speech synthesis it is desirable to have clean recorded speech that has minimal background noise. This is not easy to achieve with children as they need to feel comfortable in the recording environment and they may not be at ease in a recording studio. Another restriction that applies to collecting child speech data is the fact that the prompts are usually taken from childrens stories. It is difficult to obtain a phonetically rich corpus from this type of text. Given these challenges, the difficulty of collecting large amounts of child data could not be emphasized enough. As a result of the difficulties associated with the data collection process, only a few child voices have been developed to date. Harvesting suitable data from existing speech corpora could be a means to make more data available for child voice development. In [3], thousands of voices were built using existing speech corpora that were compiled for the purpose of training automatic speech recognition (ASR) systems. Since ASR corpora generally includes a large number of speakers, many different voices can be synthesized.

Most ASR corpora are designed to represent the phone set of the target language and should contain enough acoustic data to build models for TTS. However, recordings are more spontaneous and usually not as carefully articulated as the speech that is recorded for TTS development. ASR prompts also tend to be shorter and contain very little prosodic information. In addition, ASR corpora are usually collected with a specific application domain in mind. For instance, telephone data is recorded for telephone applications, incar data for in-car applications, etc. Only a few corpora are recorded in acoustic conditions similar to the studios where TTS data is typically recorded.

Recent studies on Hidden Markov model (HMM) based synthesis systems that use speaker-adaptive techniques (an average voice model in conjunction with model adaptation) have proven that this approach is less susceptible to imperfect data than other synthesis techniques [4], [5]. The studies demonstrated the technique's robustness to non-ideal speech data such as data recorded under unfavourable conditions and that lack phonetic balance.

This study therefore investigates the possibility to use data that has been recorded for ASR development to build a child voice for text-to-speech conversion in an HMM synthesis framework. A number of criteria were used to select ASR data. The child voices generated using ASR data were compared with a voice built with data recorded specifically for TTS in terms of intelligibility and naturalness. If the ASR data proves to be sufficient for model adaptation, more child voices could be synthesized.

This paper is organized as follows. Section II provides background to this study to place the work into context. Section III describes the corpus of child speech data and adult speech data that was used in the study. Section IV gives an overview of the criteria that were used to select data for model adaptation from the ASR child corpus. The voice evaluations and the corresponding results are described in Section V and concluding remarks are presented in Section VI.

#### II. BACKGROUND

In previous research, HMM-based synthesis using speaker adaptation was applied to generate synthetic child speech

<sup>\*</sup>This work was supported by the CSIR.

<sup>&</sup>lt;sup>1</sup>Avashna Govender and Febe de Wet are researchers at the Human Language Technology Research Group at the CSIR in Pretoria, South Africa agovenderl@csir.co.za, fdwet@csir.co.za

 $<sup>^2</sup>Bakari$  Nouhou is a Computer Science student at Enseirb-Matmeca in Bordeaux, France <code>bnouhou@enseirb.fr</code>

[1]. No South African English data was available at the time and therefore data was collected from two children during that study. The challenges that were faced in collecting this data was no different from the challenges faced by other researchers [2]. In addition, it was found that by adapting an adult gender independent initial model using Constrained Structural Maximum *a Posteriori* Linear Regression (CSMAPLR) followed by Maximum*a Posteriori* (MAP), one could generate slightly natural and intelligible speech. However, using this approach with the limited amount of data available leads to only a limited amount of voices that can be synthesized at a time.

In the study presented in [3], ASR data was used to train an initial model that was adapted to various adult target speakers. In this way thousands of adult voices were produced using HMM-based synthesis. By using this type of data, which is quite different from the data typically used to train TTS systems, the authors investigated the possibility of using non-ideal data to train TTS voices. In this way many more voices could be created than was previously possible for TTS.

In this study we decided to apply the same approach to generate child voices since there are currently no South African English child voices available. However, a corpus of South African English ASR child data also does not exist. We therefore tried to identify a data set that is as similar as possible to the target population. A corpus of American English ASR child data was subsequently chosen for an initial study. If results were promising, we could attempt to collect South African English ASR child data - which is slightly easier than collecting TTS data - for future work.

Initially, we attempted to build an average voice model using the child ASR data but the quality of the generated voices was quite bad. This was due to the lack of phonetic coverage in the data and the data from a single speaker was too little to build a good quality average voice model which is critical for the adaptation process [6]. Therefore, instead of using the ASR data as an initial model we decided to use it as adaptation data. In this way, more target speakers could be made available and more voices could be adapted.

#### **III. MATERIALS**

#### A. Child speech data

The CMU kids corpus was used in this study [7]. The corpus was recorded for the purpose of ASR development and comprises recordings of children in the age group of 6 to 11 years old. The corpus consists of 76 speakers in total of which 24 are male and 52 female.

The text that was used to develop the corpus was selected from children's magazines. It includes 356 unique sentences that corresponds to 3576 words, of which 878 are unique. The corpus contains 5180 utterances, 3434 female and 1746 male. The speech was recorded at 16 000Hz using a Sennheiser headset. The speech was directly recorded using software developed at Carnegie Mellon to record people reading aloud. The sentences were presented on the screen and when the system detected speech it would begin recording and when silence was detected the recording was automatically turned off.

Due to the difficulties experienced during recordings with children, most of the data collected is considered imperfect as the children were only asked to read each sentence once. As a consequence, the data contains many mispronunciations, whispers and various types of background noise. All mispronunciations and background noise was annotated in the transcriptions.

#### B. Adult speech data

In a previous study on child speech synthesis it was determined that it is feasible to use HMM adaptation to develop child TTS [1]. The results indicated that a mixed average voice model (AVM) trained using one adult male (bdl) and one adult female (clb) speaker from the CMU-ARCTIC database was the best initial model for adaptation to a boy child target speaker. In this study, the same AVM was used for the adaptation experiments. The CMU-ARCTIC database contains approximately 1 132 phonetically balanced sentences for each speaker [8].

#### **IV. METHODS**

The aim of this study was to select a sufficient amount of adaptation data from the ASR corpus to create various child voices using model adaptation. The  $F_0$  of speakers in this age group are so close to each other that possible differences in  $F_0$  was considered to have a minimum effect. Speaker dependence was not considered in this study as the main idea was to synthesize as many unique child voices as possible without reproducing any specific speaker's voice. Several subsets of data were selected from the corpus according to specific criteria informed by [9]. The criteria that were considered are as follows:

- 1) Clean data (with regards to transcription), no mispronunciations or mistakes in the recordings.
- 2) Data including mispronounced words.
- 3) Number of words in the utterance.
- 4) Rate of speech.
- 5) Maximum  $F_0$ .

The hypothesis was that dividing the data into clusters which share the same properties would produce better results than attempting to use all of the data at once. In other words, placing constraints on imperfect data could result in selecting data that could yield good quality voices.

#### A. Criterion 1: Clean Data

For this subset, data was selected by analyzing only the transcription files of the corpus. The corpus was transcribed differently to the conventional way of transcribing data due to the imperfections in the data. For example, the background sounds were transcribed as well as mispronunciations. A few examples of how the audio was transcribed are presented in Table I.

Only the utterances that were considered to be completely clean (without any mispronunciations or background noise) were selected for this subset. From a total of 9 hours of child

TABLE I: Conventional vs child data transcription

Conventional	Unconventional
1. the country had a war	[whisper] the country [noise] had a war [noise]
2. scientists keep learning new things about dinosaurs	scientists keep learning [noise] /TH/ new things about /D AY N AA S AO R Z/

speech data, approximately 1 hour remained after applying this criterion.

#### B. Criterion 2: Imperfect Data

For this subset, data was selected in a similar manner to the clean subset but only transcriptions that contained mispronounced words were considered. During training, the mispronounced words were mapped to an arbitrary value that would never need to be synthesized (for example "xyz") and that arbitrary value was then mapped to silence. This was an attempt to use the data without changing the labels and simply "ignore" the mispronounced segments. An example of this type of transcription is shown in Table I (Example 2).

#### C. Criterion 3: Number of correct words and rate of speech

For this subset, data was selected according to two criteria: the number of correctly pronounced words in each utterance and the rate of speech.

To create the first subset, the utterances that contained a minimum of 3 words were considered as it was assumed that 3 words is the minimum amount of words required to formulate a syntactical sentence. A few more subsets were then created by incrementally increasing the number of correctly pronounced words. These subsets included 3-5 words, 5-10 words and more than 10 words.

The speech rate was computed by calculating the number of words per second measured at utterance level. The reason why the calculation was performed at utterance level is because of how children read: their speech rate tends to fluctuate depending on their emotional engagement with what they are reading.

Combined subsets were then created by first analysing the subsets created for the first criterion (number of correctly words pronounced) and then choosing speech rates such that a sufficient amount of data could be selected.

#### D. Criterion 4: Maximum $F_0$

 $F_0$  subsets were created by clustering the data according to maximum  $F_0$  values. The corpus was then divided into two subsets based on the  $F_0$  values. The first subset included data with a maximum  $F_0$  between 180 to 300Hz. The second set contained utterances with a maximum  $F_0$  of more than 300Hz.

In addition, subsets were created by combining the data selected according to the number of correct words and rate of speech with the  $F_0$  data sets. A summary of all subsets is presented in Table II.

#### V. VOICE EVALUATION

#### A. Evaluation 1

Due to the large amount of subsets created, it was not feasible to evaluate all 12 output voices subjectively. Therefore, the first round of evaluations were performed internally by two listeners. The evaluation consisted of an intelligibility test where each sample was transcribed three times using a set of test data. The same sentences were listened to for each voice<sup>1</sup>. A word error rate (WER) was calculated for each transcription and the voices were ranked accordingly. The results for the voices that yielded WERs below 50% are shown in Figure 1.



Fig. 1: Intelligibility Test WER

The following observations were made from these results: When using the subsets of clean and imperfect data separately, they performed worse than when they were used together. This could be due to the fact that using more data yields better quality output. However, the aim of this study is to show that one should carefully consider what data is combined rather than using data combinations blindly. The voices defined from Subsets 12 and 7 resulted in the lowest WERs. Subsets 12 and 7 share some of the data in terms of the number of words and speech rate.

Given these results, additional subsets were created from Subsets 7 and 12. The subsets isolated the individual criteria to clarify which criterion or combination of criteria resulted in the most intelligible voices. Two additional subsets were created by randomly selecting data from the corpus. This was to test whether the data selection criteria yielded the positive results and not just random data selection. The new subsets are summarized in Table III. The listening

<sup>&</sup>lt;sup>1</sup>Since only two listeners performed the test at this stage, listening to the same sentence could have influenced the results as listeners might have gained prior knowledge and transcribed based on the prior knowledge rather than what is really said in the sample. This risk was mitigated by randomizing the sequence in which the utterances were presented to the listeners.

TABLE II: Summary of subsets

Subset	Criteria	No. of Utts	Duration (hr:min)
Subset 1	Clean transcriptions	836	01:12
Subset 2	Transcriptions with mispronounced words	3019	07:43
Subset 3	Combination of Subset 1 and 2	3855	08:55
Subset 4	More than 3 words, speech rate between 0.5 and 1.5	2050	03:46
Subset 5	3 to 5 words, speech rate between 0.5 and 1.5	1274	01:48
Subset 6	5 to 10 words, speech rate $> 1.5$	1218	01:41
Subset 7	More than 10 words, speech rate $> 1$	704	01:22
Subset 8	More than 10 words, speech rate $< 1$	88	00:21
Subset 9	maximum $F_0$ between 180 to 300Hz	990	01:28
Subset 10	maximum $F_0$ between > 300Hz	2710	04:32
Subset 11	More than 5 words, speech rate $> 1.5$ , maximum $F_0$ more than 300Hz	918	00:07
Subset 12	More than 5 words, speech rate $> 1.5$ , maximum $F_0$ less than 260Hz	107	01:13

test performed previously by two internal candidates was subsequently repeated using the same sentence set as before. The WERs calculated are shown in Figure 2.



Fig. 2: Intelligibility Test WER for voices built from additional data sets

The following observations were made from these results: Subset 13 (More than 10 words), performed the best. This result seems to be related to the results observed for Subsets 7 and 12 in the first round. Both these sets have intersections with Subset 13. Subset 17 performed second best with an additional criterion of speech rate greater than 1.5. Adding this limitation to the data resulted in a 1% decrease in WER.

Following this trend, by adding the  $F_0$  criterion, a further deterioration in the results were observed as shown by Subset 19, based on the combination of all 3 criteria. However, it is noted that this particular subset had only 13 utterances with a duration of a minute which is usually insufficient data to perform adaptation. This could explain why it performed so poorly. Subset 18 also performed well but it is not clear as to which properties of the data its performance can be ascribed. The randomly selected subsets did not perform well and the possibility that the improved intelligibility of the voices is due to random properties of the data can therefore be rejected.

If the three best voices from this evaluation were selected for subjective evaluation, it would be purely based on intelligibility and not at all on naturalness. Therefore another test was performed by just listening to the voices that had a WER less than 50%. In this manner, five voices were selected and ranked in terms of naturalness, as shown in Table IV.

TABLE IV: Naturalness Ranking

Rank	Voice
Rank 1	Subset 17
Rank 2	Subset 7
Rank 3	Subset 13
Rank 4	Subset 12
Rank 5	Subset 11

Comparing these results with the intelligibility results, Subsets 17, 13 and 7 were selected and evaluated in the next round of evaluations. Selecting Subset 17 and 13 was an easy choice as they performed well for both the naturalness and intelligibility. Since Subset 7 and 12 performed better than Subset 11 in both tests, the choice was between Subset 7 and 12. Finally, Subset 7 was chosen because it seems to be consistent in terms of criteria with the other two identified subsets.

#### B. Evaluation 2

For this round of evaluations, a formal listening test was conducted using a web-based evaluation test. 20 listeners participated in the test. The listening test consisted of two sections:

- 1) Mean-Opinion-Score (MOS) test to evaluate naturalness
- 2) Transcription test to evaluate intelligibility

For each test four voices were evaluated randomly. Three voices were selected during Evaluation 1 and the fourth voice was a synthesized voice of a male target child speaker as described in [1]. In this way, the voice trained using data recorded specifically for speech synthesis could be compared with the voices trained using the ASR data.

For the MOS test the listeners were asked to listen to an utterance and rate the voice on a 5-point scale where the digits corresponded to the following: 1 - Completely Unnatural, 2 - Unnatural, 3 - Slightly Natural, 4 - Natural and 5 - Completely Natural.

For the transcription test, a test set of 24 utterances were synthesized. Each listener listened to and transcribed three utterances per voice. Special care was taken to avoid a single listener evaluating the same sentence twice. The listeners first transcribed 12 sentences that were semantically predictable (SPS) and then a second set of 12 sentences that were

TABLE III: Additional data sets

Subset	Criteria	No. of Utts	Duration (hr:min)
Subset 13	More than 10 words	508	01:09
Subset 14	Speech rate $> 1$	2668	03:48
Subset 15	More than 10 words, speech rate $> 1$ , maximum $F_0$ less than 260Hz	20	00:02
Subset 16	More than 5 words, speech rate $> 1$	2095	03:17
Subset 17	More than 10 words, speech rate $> 1.5$	355	00:39
Subset 18	More than 5 words and speech rate > 1, maximum $F_0 < 260$ Hz	105	00:10
Subset 19	More than 10 words, speech rate > 1.5, maximum $F_0 < 260$ Hz	13	00:01
Subset 20	Randomly selected utterances	107	00:10
Subset 21	Randomly selected utterances	107	00:10

semantically unpredictable (SUS). The results of these tests are presented in Figures 3 and 4.



Fig. 3: Subjective MOS scores

The MOS scores show that both the TTS voice and the Subset 13 voice have a median of 2 which is considered to be unnatural. Subset 7 also has a median of 2 but the boxplot is positively scewed. This means that most of the scores lie between 2 and 3 where 3 is considered to be slightly natural. Subset 17 has a median of 2.5 and is also positively scewed. Overall, Subset 17 is the most natural of the 4 voices evaluated, but still requires improvement as the naturalness is only considered to be slightly natural. The limited amount of adaptation data used could be the reason behind the poor naturalness of the voices.

The WERs in Figure 4 indicate that for the SPS test, the TTS voice, as expected, has the lowest WER of 25% followed by Subset 7 and Subset 17 with WERs of 46% and 47% respectively. Subset 13 has the highest WER of 56% which is a gap of 10% from the other two ASR subsets. For the SUS test, the TTS voice again had the lowest WER of 58%, followed by Subset 17 with 68% WER. The remaining subsets had exceptionally high WERs of 79% and 84% respectively. It is evident that the WER increased substantially when using unpredictable sentences instead of predictable sentences. However, the results seem to be consistent between the two sets of results.

Overall, Subset 7 and Subset 17 are comparable to the TTS voice for both naturalness and intelligibility. Subsets 7

Intelligibility Test Results





Fig. 4: Intelligibility Test WER for SPS and SUS

and 17 have an intersection in terms of the number of words in the utterance as well as the the rate of speech. That is, the data selected for those subsets have utterances of more than 10 words and a speech rate of more one word per second. However, the specific subsets that performed the best in this study may not necessarily be relevant to another corpus. On the other hand, this work shows that by carefully selecting data based on specific criteria relevant to the corpus, one can narrow down the data in such a way as to select only the data that could generate good quality synthesized voices, given that the data selected is imperfect.

#### VI. CONCLUSION

The results of this study indicate that, in terms of naturalness, the voices trained with ASR data is comparable to a voice trained using TTS data. In fact, one of the voices performed better than the TTS voice. However, there is room for improvement. In terms of intelligibility, the WERs were not as close to the TTS voice but it indicates that, with proper selection methods, ASR data could be used as an alternative to TTS data when only a small amount of data is available for TTS model development. Although it can be argued that these results could be dependent on the type of data that is contained in the database, it can also be said that by carefully clustering a sufficient amount of data according to specific criteria, comparable voices can be synthesized. In future work, a more in-depth analysis of the data selection criteria will be conducted. In doing so, a more accurate conclusion can be drawn as to why certain criteria yield better results than others. In addition, data needs to be

selected such that the subsets can be of equal size in terms of the duration, as it was evident that subsets with very few utterances (duration of less than 30 minutes) performed quite poorly in the evaluations.

#### References

- A. Govender, F. de Wet, and J. R. Tapamo, "HMM adaptation for child speech synthesis," in *INTERSPEECH 2015*, (Dresden,Germany), pp. 1640–1644, Sept.2015.
- [2] O. Watts, J. Yamagishi, K. Berkling, and S. King, "HMM-based synthesis of child speech," in *Proc. of The 1st Workshop on Child Computer and Interaction.*, (Crete,Greece), Oct.2008.
- [3] J. Yamagishi, B. Usabaev, S. King, O. Watts, J. Dines, J. Tian, Y. Guan, R. Hu, K. Oura, Y.-J. Wu, *et al.*, "Thousands of voices for HMMbased speech synthesis–Analysis and application of TTS systems built on various ASR corpora," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 18, no. 5, pp. 984–1004, 2010.
- [4] J. Yamagishi, T. Nose, H. Zen, Z.-H. Ling, T. Toda, K. Tokuda, S. King, and S. Renals, "Robust speaker-adaptive HMM-based textto-speech synthesis," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 17, no. 6, pp. 1208–1230, 2009.
- [5] J. Yamagishi, H. Zen, Y.-J. Wu, T. Toda, and K. Tokuda, "The HTS-2008 system: Yet another evaluation of the speaker-adaptive HMMbased speech synthesis system in the 2008 Blizzard Challenge," 2008.
- [6] J. Yamagishi, Average-voice-based speech synthesis. PhD thesis, Tokyo Institute of Technology, 2006.
- [7] University of Pennysylvania, "Linguistic data consortium CMU kids corpus." http://catalog.ldc.upenn.edu/LDC97S63. Date Accessed:September 2015.
- [8] J. Kominek and A. W. Black, "The CMU arctic speech databases," in Fifth ISCA Workshop on Speech Synthesis, 2004.
- [9] E. Cooper, "Data selection for Text-to-speech for Low-Resource languages," 2015. Poster presented at Summer School on Speech processing, July 27–31, University of Crete, Crete, Greece.

2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech) Port Elizabeth, South Africa, November 26-27, 2015

# Autonomous Prediction of Performance-based Standards for Heavy Vehicles

Robert Berman1, Richardt Benade1, Benjamin Rosman1,2 1 CSIR, Pretoria, South Africa 2 School of Computer Science and Applied Mathematics, University of the Witwatersrand

rberman@csir.co.za, rbenade@csir.co.za, brosman@csir.co.za

Abstract-In most countries throughout the world, heavy vehicle use on public roads are governed by prescriptive rules, typically by imposing stringent mass and dimension limits in an attempt to control vehicle safety. A recent alternative framework is a performance-based standards approach which specifies onroad vehicle performance measures. One such standard is the low-speed swept path, which is a measure of road width required by a vehicle to complete a prescribed turning manoeuvre. This is typically determined by physical testing or detailed vehicle simulations, both of which are costly and time consuming processes. This paper presents a data driven, detailed model to predict the low-speed performance of an articulated vehicle, given only the vehicle geometry. The development of a lightweight tool to predict the swept path of an articulated heavy vehicle, without the need for detailed simulation or testing, is discussed.

Keywords— Performance-based standards; Vehicle safety; Heavy vehicle performance; Regression; Support vector machines

#### I. INTRODUCTION

In most countries throughout the world, heavy vehicle use on public roads is governed by prescriptive rules. In South Africa, the National Road Traffic Act (NRTA) specifies legal mass and dimension limits for all vehicles that operate on public roads. Two of the main constraints placed on heavy vehicles are overall length and mass limits, specified as a maximum of 22 m and 56 000 kg, respectively. These prescriptive limits are both easy to understand as well as enforce, however they do not inherently regulate vehicle safety as factors that influence the actual on-road performance of the vehicle are largely not governed.

A recent alternative framework is the performance-based standards (PBS) approach which specifies actual on-road vehicle performance measures, as opposed to merely limiting what the vehicle looks like. The South African PBS pilot project has been operational since 2007 (CSIR Built Envoronment, July 2015) and currently boasts 160 so-called Smart Trucks operating on designated routes across the country. A typical 9-axle, 73 tonne, 26 m Smart Truck, known as a B-double (due to the roll-coupled trailers, the towing mechanism and the two trailers in addition to the tuck tractor) is shown in Figure 1.

For a PBS Smart Truck to be granted a permit to operate on the road, a detailed vehicle safety assessment of the vehicle is required, considering all aspects of the vehicle: engine,



Fig. 1. Representation of a PBS B-Double

suspension, springs, dampers, tyres, vehicle mass properties and payload mass properties. These safety assessments require input from the truck tractor original equipment manufacturer (OEM), the trailer OEM, as well as the vehicle operator, resulting in the PBS assessment being a costly and time consuming process due to many hours of costly physical and software testing. The vehicle safety assessment process is also highly iterative, with the OEMs creating a vehicle layout design, the assessment is then conducted and the results forwarded to the OEM to update the design if necessary. There are currently no simple, standalone tools available for the OEMs to calculate or predict the vehicle performance without conducting a formal assessment. Therefore, a simple model that is able to predict vehicle performance given simple geometric vehicle properties will provide great insight for the OEMs as well as reduce the time and financial costs of the formal vehicle safety assessment.

To date, vehicle safety assessments have been conducted using multi-body dynamic analysis software simulation packages, requiring a detailed and comprehensive understanding of the physics and dynamics of the entire vehicle system as well as its subsystems such as suspension and tyres. This paper presents a data driven approach to predict the lowspeed performance of articulated heavy vehicles, requiring only simple geometric vehicle properties and no knowledge of the mechanics of the system. The model presented in this paper utilises commonly used supervised machine learning methods.

The remainder of this paper is structured as follows. In Section II we provide a detailed discussion of the performance standards against which these vehicles are measured. In Section III, a brief background to the learning techniques is given. The overall structure of the system is presented in Section IV, with Section V and Section VI covering the system performance and other related work respectively. Finally, we present a discussion of the data driven model and the conclusions of the research.

#### II. THE PERFORMANCE STANDARDS

There are two main components of a PBS vehicle safety assessment, namely low-speed directional standards and high-speed stability standards. Each standard is given a performance level, either pass/fail or level 1 to level 4, with level 1 having the most stringent performance criteria. This paper focuses on the five low-speed standards, which, for a specific vehicle, are measured from a prescribed 12.5 m radius, 90° degree turn. The five low-speed standards are defined as (NTC, 2008):

**Low-Speed Swept Path (LSSP)** is the amount of road width required by the vehicle when executing the prescribed low-speed 90° turn, as the trailing units track inside of the path followed by the hauling unit. LSSP is given a performance rating of level 1, level 2, level 3, level 4 or fail.

**Frontal Swing (FS)** is the amount that the front outside corner of the hauling unit swings beyond of the exit tangent of the widest section of the vehicle at the completion of the low-speed 90° turn. FS has a pass/fail performance level.

**Difference of Maxima (DoM)** and **Maximum of Difference (MoD)** pertain to the amount by which the front outside corner of a semitrailer swings out beyond that of the path of the hauling unit or preceding semitrailer. DoM and MoD both have pass/fail performance levels.

**Tail Swing (TS)** is the amount which the rear outside corner of a vehicle unit swings out at the commencement of the prescribed low-speed 90° turn. This may cause collisions with objects in adjacent lanes or on the roadside. TS is given a performance rating of level 1, level 2, level 4 or fail.

Figure 2 shows the critical points of a vehicle during the prescribed turn for a) LSSP and b) MoD and DoM.



#### **III. BACKGROUND TO LEARNING TECHNIQUES**

The model presented in this paper uses commonly implemented machine learning techniques and methods to predict the low-speed performance of heavy vehicles.

Due to the availability of vehicle geometrical data as well as ground truth outputs, supervised learning techniques utilised for this study. The requirement for an appropriate technique was to accurately map the vehicle geometric parameters to low-speed performance.

The first supervised machine earning technique used in the model is a multilayer perceptron (MLP). MLP networks have been in use for many years, but have made a recent resurgence in deep learning, as evidenced by (Raiko, Valpola, & LeCun., 2012). They note that this rise has followed the invention of unsupervised pretraining, however there has also been a modern trend to utilise tradational back-propagation as this method is abable of giving sufficent accuracy.

The second supervised machine learning technique used in the prediction model is support-vector machines (SVM). SVMs have been used for a wide variety of learning problem for classification, regression as well as other learning tasks. The LIBSVM library has been used extensively and is widely cited in the literature (Chang & Lin, 2011). SVMs represent a powerful technique for general non-linear classification (Meyer & Wien, 2014), and are used for such in this study.

#### IV. SYSTEM ARCHITECTURE

The mechanics of articulated vehicles executing turning manoeuvres are such that the geometric input parameters for the low-speed model are not uncoupled, and cannot be considered to be independent (Winkler & Aurell, 1998). Winkler and Aurell also show that the mechanics of articulated vehicle turning are non-linear and non-elementary.

In the next section, a model comparing the geometric inputs directly to the overall vehicle level is presented, however this overly naïve model is not able to capture any of the dependencies of these variables, nor is it able to give any insight into the relative performance of the vehicle in each standard. In order to capture the input parameter dependencies, as well as the system non-linearity, a complex data driven model was created to predict the level of each low-speed standard.

In practice, the exact LSSP value of an articulated vehicle directly affects its ability to navigate a given intersection. The magnitude of the LSSP provides the OEM with a greater insight into the vehicle performance than that of the LSSP level, and as such, a regression model was selected for this standard. The output of the numeric regression is fed through a thresholding operation to give the LSSP level.

The remainder of the standards relate more to overall vehicle safety whilst the vehicle navigates the turn than they do to its ability to navigate the turn. Thus, for these standards, the level is of greater interest than the magnitude of the output, and as such, classification models were implemented to directly output the level. The model for each standard is presented as a system with the output from each sub-system combined in a final layer due to prior knowledge of the definition of each

Fig. 2. a) LSSP, b) FS, MoD and DoM

standard, as such, it would be redundant to learn these. This represented schematically in Figure 3.



Fig. 3. System Block Diagram

A diagram showing the detailed system architecture, illustrating the relationship of all the low-speed standard models to the overall output level, O, is given in Figure 4.



Fig. 4. System Diagram

In addition to capturing the nuances of each standard, the system presented here provides greater transparency with the level of each standard, which will aid the justification to the external legislative bodies for implementing such a system. This system will also provide additional benefits to the OEMs during the design of articulated vehicles, as mentioned above.

A simplified system model was also created, relating the 22 input parameters directly to the overall low-speed level. This simple model gives insight into the overall performance of the vehicle combination, but is not able to give any insight into the relative performance across each standard. This simple model will provide a useful check to OEMs to confirm whether or not their design meets the required level. The detailed model however, will be a more valuable tool, as it can be integrated into the design process to not only confirm a design's performance, but also to optimise vehicle designs.

The complex system presented here will also form the basis of a model that also includes the high-speed standards. This modular model easily allows for the addition of individual standards, ensuring that the detailed insights into the performance for each standard are retained, without the need to recreate the entire model.

#### V. DATA ANALYSIS

A set of 10 000 simulations, with randomly selected vehicle geometrical parameters for a B-double were run to obtain the ground truth values for all five low-speed standards. In total, 22 parameters were selected as inputs, seven for each of the first two units and eight for last unit. These data were then used to create a regression model for LSSP and classification models for each remaining of the four standards. These models were then combined into a single system, given above, to predict the level for each standard, as well as the overall low-speed level for the combination. The regression and classification models were created in Weka 3.6, using default parameters, unless otherwise stated.

The MLP's used in this study all utilised sigmoid activation functions. The number of hidden units and nodes were selected according to the combination that yielded the greatest accuracy, starting with a single hidden layer and the number of nodes equal to half the sum of the number of inputs and outputs.

The regression model for LSSP utilised a MLP comprising two hidden layers, with six and three nodes respectively. The hierarchical structure of the simple MLP was capable of capturing the complexity of the LSSP mechanics without being overly complex itself. The number of basis functions was selected in advance, ensuring the simplicity of the model, as well as ensuring a compact model that is able to quickly process new data (Bishop, 2006). The classification models for FS, DoM and MoD also utilised MLP for the same reasons. It was found that for these standards, the MLPs were able to achieve good accuracy for classification, whilst limiting the number of false positives, which are highly undesirable.

The MLP classification of the FS standard comprised three hidden layers with twelve, ten and five nodes, respectively. The MLP for DoM contained only a single hidden layer with three nodes, while that of the MoD standard comprised four hidden layers with 21, fifteen, ten and five nodes respectively.

The classification model for TS utilised SVM for the four class classification. Multiclass SVM models have undergone a number of development iterations to improve their capability of fundamentally being a two-class classifier (Bishop, 2006). The LibSVM library in Weka, with a radial basis function (RBF) kernel, was selected to capture any non-linearity in the TS data (Witten & Frank, 2005). A cross-validation and grid-search, was conducted to select the two RBF kernel parameters, (C,  $\gamma$ ). The parameters which gave the greatest accuracy were a gamma of 0.017 and a cost of 48.

The accuracy and performance of all the low-speed models are given in Table 1, with the confusion matrices for the four classification models given in Table 2 to Table 5.

TABLE I. PERFORMANCE OF LOW-SPEED PREDICTION MODELS

Std	Alg	Corr. Coeff.	Relative Absolute Error (%)	Relative Squared Error (%)	Classified Correctly (%)	False Positive (%)
LSSP	MLPr	0.9995	2.8615	3.0176		
TS	SVM <sub>c</sub>				96.31	
FS	MLP <sup>c</sup>				95.72	7.83
DoM	MLP <sub>c</sub>				91.14	6.12
MoD	MLP <sub>c</sub>				98.3	0.84
Combined					95.36	3.29
Direct	MLP <sub>c</sub>				94.67	1.76

TABLE II	TS CONFUSION MATRIX
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1	

Level 1	Level 2	Level 4	Fail	
9043	26	41	1	Level 1
65	6	45	6	Level 2
56	22	158	22	Level 4
3	3	47	3	Fail

TABLE III.	FS CONFUSION MATRIX		
Level 1	Fail	_	
5968	173	Level 1	
198	2330	Fail	

TABLE IV.	DOM CONFUSION MATRIX		
Level 1	Fail		
1366	259	Level 1	
257	3944	Fail	

TABLE V.	MOD CONFUSION MATRIX		
Level 1	Fail	_	
2125	80	Level 1	
48	5651	Fail	

The performance of the simple overall model relating inputs to overall vehicle level is given in Table 6.

TABLE VI SIMPLE OVERALL LEVEL CONFUSION MATRIX

Level 1	Level 2	Level 3	Level 4	Fail	_
1055	6	0	0	184	Level 1
11	463	6	1	76	Level 2
0	5	341	0	34	Level 3
17	0	2	5	55	Level 4
87	33	16	0	7603	Fail

#### VI. RELATED WORK

The majority of simplified models for predicting PBS performance of heavy vehicles comprise simplified models of the system mechanics and physics, notably the static rollover hreshold (SRT) calculator, that has been written into egislation in New Zealand (de Pont, Baas, Hutchinson, & Kalasih, 2002). This tool takes simplified overall vehicle parameters as inputs and calculates SRT from first principles.

A similar first principle approach is the complex vehicle nodel for turning that uses the physics and mechanics of the vehicle to calculate the trajectory of a vehicle in a turn Winkler & Aurell, 1998). This model utilises a comprehensive system of equations to model the physics of the vehicle and the nanoeuvre. It gives highly accurate results, yet requires a full understanding of the physics of the system in order to use and successfully implement.

Dessein et al. presented a simplified, third order polynomial, regression model to estimate LSSP (Dessein, Kienhofer, & Nordengen, 2012). This model gives good accuracy for a generic articulated vehicle, but is unable to give any insight into the other four low-speed PBS standards.

De Pont presented a pro-active approach to ensure LSSP compliance through a so called pro-forma approach (De Pont, 2010). The pro-forma design specifies limits on the geometric properties of an articulated heavy vehicle such that the required LSSP level is met. Benade et al. expanded on this approach with a pro-forma design to additionally include the FS and TS standards (Benade, Berman, Kienhofer, & Mordengen, 2015). This adapted pro-forma design is limited by the allowed range for each input parameter and thus is only applicable to a narrow spectrum of vehicles.

#### VII. DISCUSSION AND CONCLUSION

The existing tools that are used to calculate the low-speed performance of articulated heavy vehicle have been shown to give good accuracy, however require a detailed understanding of vehicle dynamics to successfully implement. The development of pro-forma designs has sought to introduce a data driven approach to vehicle design, but to date have been only partially successful, and have been limited to stringent constraints for a specific vehicle combination.

The data driven model presented in this paper provides an accurate tool that is naïve towards the physics of the system, yet is able to capture the nuances of the physics. The system gives insights into the individual low-speed standards as well as overall vehicle performance based on basic geometric vehicle properties.

This model is not limited by vehicle layout and is able to generalise to any B-double combination. Current work includes an extension of the low-speed prediction model to include a wider variety of vehicle configurations, with differencing numbers of trailers to improve the applicability of this model to heavy vehicle industry in South Africa. Future work will include further expanding the data driven model to predict vehicle performance in the high-speed stability standards.

#### **ACKNOWLEDGEMENTS**

The authors gratefully acknowledge the helpful comments of the anonymous reviewers.

#### REFERENCES

- [1] CSIR Built Envoronment, "Smart Truck Programme Rules for The Development and Operation of Smart Trucks as Part of The Performance-Based Standards Research Programme in South Africa ver 17", Pretoria: CSIR Built Envoronment, July 2015.
- [2] NTC, "Performance based standards scheme The standards and vehicle assessment", Melbourne, Australia: National Transport Commission, 2008.
- [3] Raiko, Tapani, Harri Valpola, and Yann LeCun. "Deep learning made easier by linear transformations in perceptrons." International Conference on Artificial Intelligence and Statistics, 2012.
- [4] Chang, Chih-Chung, and Chih-Jen Lin. "LIBSVM: A library for support vector machines." ACM Transactions on Intelligent Systems and Technology (TIST) 2.3, 2011.
- [5] Meyer, David, and FH Technikum Wien. "Support vector machines." The Interface to libsvm in package e1071, 2014.

- [6] C. B. Winkler, J. Aurell, "Analysis and testing of the steady-state turning of multiaxle trucks", 5th International Symposium on Heavy Vehicle Weights and Dimensions, 1998, pp. 135-161.
- [7] C. M. Bishop, "Pattern Recognition and Machine Learning", New York: Springer, 2006.
- [8] I. H. Witten, E. Frank, "Data Mining: Practical Machine Learning Tools and Techniques", San francisco: Elsevier, 2005.
- [9] J. J. de Pont, P. Baas, D. Hutchinson, and D. Kalasih, "Including Performance Measures in Dimensions and Mass Regulations", 7th International Symposium on Heavy Vehicle Weights and Dimensions, vol. 41001, 2002, pp 349-358.
- [10] T. Dessein, F. Kienhofer, and P. Nordengen, "Determining the Optimal Performance Based Standards Heavy Vehicle Design", 12th International Symposium on Heavy Vehicle Transportation Technology. Stockholm: International Forum for Road Transport Technology, 2012.
- [11] J. J. de Pont, "The development of pro-forma over-dimension vehicle parameters", Auckland, New Zealand: TERNZ, 2010.
- [12] R. Benade, R. Berman, F. Kienhofer, and P. Nordengen, "A Pro-Forma Design for Car-Carriers: Low-Speed Performance-Based Standards", 34th South African Transport Conference, Pretoria, 2015, pp. 253-265.

# Rapid Prototyping of Small Wind Turbine Blades Using Additive Manufacturing

Sean Poole<sup>1</sup> and Russell Phillips<sup>2</sup>

Abstract—For practical research results at NMMU on small wind turbine blade design, a manufacturing method was required to rapid prototype various iterations of blade designs. The requirements of the manufacturing method were a rapid (a few days), cost effective and accurate technique to produce full size small wind turbine blades (up to 1.4 m long) that could withstand short term testing. From the manufactured wind turbine system, results could be obtained and design changes made to the following iteration of the blade design. The process used to manufacture the blades involves five steps:

- 1) A program to design the turbine blade.
- 2) A program to develop a CAD model from the designed blade parameters.
- 3) A CAM programme to process the CAD for 3D printing.
- 4) 3D printing (additive manufacturing) of the turbine blade geometry using FDM (Fused Deposition Modeling).
- 5) Reinforcement of the PLA plastic printed blade.

For this research, the 3D printer has been designed and manufactured. The first few iterations of blade designs have been manufactured and the testing of the blades is currently under way. Future testing of the long term structural characteristics of turbine blades produced using this method is recommended and planned. This documented research is drawn from a section of Poole's PhD research (*Optimisation of a Mini Horizontal Axis Wind Turbine to Increase Energy Yield During Short Duration Wind Variations*) [1].

#### I. INTRODUCTION

This aim of this paper is to document the developed method for a fast turn around time manufacturing method for the short term testing of small wind turbine blades. The rationale for a rapid prototyping method was to allow for various iterations of blade designs to be produced quickly and with minimal effort, in order to encourage iterative progress without the stifling factors of substantial effort, time, and costs (as with the manufacturing process of a conventional moulded composite blade). An additive manufacturing technology process (FDM) was chosen due to its geometry flexibility, sufficient accuracy, rapid prototyping abilities, and low manufacturing costs. The rest of this paper provides details of the following processes involved with the chosen manufacturing technique and is graphically represented by Figure 1:

1) A Labview program was written to design a turbine blade based on BEM (Blade Element Momentum) theory [2] and the inputs of the user. The parameters of this blade were then exported for step 2.

- 2) Once the turbine blade parameters have been determined and calculated (aerofoil profiles, pitch, and chord length), an accurate CAD model of the blade is required. Using the CAD software Rhinoceros 3D [6] and the add-in package Grasshopper [5], a programme was written to import the blade parameters and automatically model the turbine blade with a few adjustable parameters (hub mounting dimensions, tip parameters, etc).
- 3) From the CAD model developed in Step 2, the blade model is processed using Repetier-Host [3] to create the machine code required to additive manufacture (3D print) the full-scale shell of the turbine blade.
- 4) The blade is then 3D printed using PLA plastic and a FDM process.
- 5) The 3D printed shell is then reinforced using pultruded glass-fibre rods, epoxy resin, and/or micro-balloons. The mounting holes are drilled through the 3D printed drill dimples in the model. The blade is then ready for mounting and testing.



Fig. 1: Manufacturing Process Flow Diagram

#### **II. MANUFACTURING PROCESS**

The following 5 sections (as briefly described in the *Introduction*) provide the documented details of the manufacturing process for a small wind turbine blade.

<sup>&</sup>lt;sup>1</sup>Sean Poole is with Faculty of Mechanical Engineering, Renewable Energy Research Group (RERG), Nelson Mandela Metropolitan University, Port Elizabeth, South Africa seanp@nmmu.ac.za

<sup>&</sup>lt;sup>2</sup>Russell Phillips is with Faculty of Mechanical Engineering, Renewable Energy Research Group (RERG), Nelson Mandela Metropolitan University, Port Elizabeth, South Africa russell.phillips@nmmu.ac.za

#### A. BLADE DESIGN PROCESS

A Labview programme was written to design the parameters for the turbine blades based on a few user inputs and BEM theory and including Prantl tip-loss correction factor (applied to the induction factors to compensate for tip losses) [4]. The required user inputs are shown in Figure 2 as per the GUI (Graphical User Interface) of the written programme. The imported *lift* and *drag* files contain the lift and drag coefficients for the single aerofoil profile at various AoA (Angle of Attack) and for various Reynolds numbers. The inputted Coordinate file is non-critical and is used for graphical representation of the designed blade system. For a typical lift-to-drag design the AoA Offset is set to 0 and therefore the BEM process is optimized by maximizing the lift-to-drag ratio. Other input parameters include Number of Blades, Design Wind Speed, Turbine Radius, Hub Radius (which specifies the inner starting radius for the designed blade section), TSR (Tip Speed Ratio which specifies a relative turbine speed), Air Density, Number of Elements (which specifies the number of equal section divisions along the length of the blade which each represent a 2D aerofoil and airflow interaction in the BEM process), and then the Error Tolerance (which specifies the maximum allowable relative convergence error between the induction factor forces and the 2D aerofoil forces, in order for the BEM theory to be valid [4]).



Fig. 2: Example Labview User Input for the Blade Design Parameters

Once the Labview programme has run, the parameters for the designed blade (radius, chord, pitch) are exported to a file while also being represented graphically in the GUI as shown in Figure 3, and Figure 4. Figure 5 shows the power coefficient along the length of the blade and is a good representation of the theoretical efficiency relative to the theoretical maximum Betz limit of 0.59 [4].

Finally, within the Labview programme environment, a quick graphical representation of the designed system can be seen and is shown in Figure 6 and Figure 7.



Fig. 3: Example Labview Blade Design Output of Chord vs Radius



Fig. 4: Example Labview Blade Design Output of Pitch vs Radius



Fig. 5: Example Labview Blade Design Output of Local Power Coefficient vs Radius



Fig. 6: Example Labview Graphical representation of Blade Design



Fig. 7: Example Labview Graphical representation of Blade System

#### B. CAD MODEL DEVELOPMENT PROCESS

Once the turbine blade parameters have been determined, an accurate CAD model (including blade tip, and hub mounts) is required for the manufacturing process. A programme was written in Grasshopper [5] (plug-in for Rhinoceros 3D [6]) which could import the blade parameters, process the data, and then sketch the framework for the CAD model. Various parameters are adjustable in the programme which allows for tweaking of the model to suit the hub mounting requirements and blade tip requirements. Figure 8 shows the blade-hub transition control parameters.



Fig. 8: Example Grasshopper Input Parameters

The programme also does leading and trailing edge

smoothing to compensate for discrete element sectioning and discrete look-up table values (from the BEM calculated outputs). The Grasshopper output is represented by the example shown in Figure 9. The red framework shows the output of the Grasshopper programme which forms the basis of the CAD model. Once the surfaces of the model are manually created from the framework, this model can then be saved as a STL file and is ready for 3D printing CAM (Computer Aided Manufacturing).



Fig. 9: Example Grasshopper Output into Rhinoceros 3D

#### C. CAM PROCESS

The processing stage to convert the CAD model into machine code for the 3D printer (CAM) is done using a freeware software called Repetier-Host [3]. Some of the features of this software include *spiral print* and *extruded plastic width control*. These features allow the turbine blades to be printed hollow in one continuous movement spiralling upwards while traversing the outer perimeter of the blade. The hollow plastic blade would then function as a "mould" for the reinforcement casting. The printing can also be done using conventional techniques with an internal lattice structure which helps to prevent shell warpage but also dramatically increases the required printing time.

#### D. 3D PRINTING MANUFACTURING PROCESS

Since the size of the turbine blades is unconventionally large for a 3D printer, a custom 3D printer was designed and manufactured. The design was based on a Prusa I3 3D printer [7] using Marlin [8] open source software. Figure 10a shows a rendering of the CAD model of the designed 3D printer, and Figure 10b shows the realised product. All blades were printed using PLA plastic instead of ABS due to





(a) CAD of 3D Printer

(b) Realized 3D Printer

Fig. 10: Large 3D Printer

the better thermal contraction properties of PLA (less thermal contraction).

#### E. BLADE REINFORCEMENT PROCESS

The second step in the manufacturing process is to reinforce the printed blade. Various reinforcement methods were used to add structural strength to the plastic (PLA) printed blades and are listed:

- 1) **POUR FILLED** The blade was printed completely hollow. An epoxy-glass micro balloon mix was cast inside the print (Figure 11).
- SHORT FIBRE INFUSED The blade was printed completely hollow. The blade was filled with a loose short strand fibre-glass (12.7 mm) and then infused with epoxy resin (Figure 12).
- 3) **PULTRUDED ROD REINFORCED** The blade was printed with a lattice internal fill structure. Composite rods were inserted in between the lattice structure from blade root to blade tip. The tip and root of the blade were internally cast with epoxy to tie the printed blade together with the composite rods (Figure 13).

#### III. MANUFACTURING PROCESS COMPARISON

For the manufacturing processes as previously listed in Section II-E, the following are the advantages and disadvantages noted for each process:

#### A. Pour Filled

1) Advantages:



(b) Cross Section Material Diagram









- 1) Due to the hollow printed shell, only a single pass of the printer head is required along the perimeter of the blade shape. This makes continuous spiral printing possible, which speeds up the printing process. Also, the absence of of an internal structure speeds up the printing process. The time required to print a 1 m blade is between 6 and 12 hours depending on blade volume).
- 2) The hollow printed shell allows print material and costs to be saved since no internal structure is required.
- 3) Using a variable print speed with the spiral print method allows for a stronger print since relapse pass rate can remain constant independent of the perimeter length. This allows for a stronger more consistent print.
- 4) The reinforcement method is quick and simple, only requiring the mixture of the epoxy and glass balloons to be poured into the shell (and be also be kept cool to prevent exo-thermal runaway).
- 5) Once the reinforcement is complete the PLA plastic shell can either be left in situ or removed.
- 2) Disadvantages:
- Large flatter surfaces on the part geometry can undergo warping during printing. This only occurred on the high pressure side of the aerofoil profile and was therefor considered not a major concern.
- 2) Due to the large amount of epoxy used, the risk for exo-therm runaway was high and water cooling was required to mitigate this risk.
- 3) Cracking or leaking of the shell during water cooling



(a) Manufactured Blade

(b) Cross Section Material Diagram



risked the inclusion of water inside the blade from the water cooling process.

- 4) This method was the weakest of the processes tested.
- 5) Unbalancing between blades weights is possible during this process as blade volume (and therefore weight) can change slightly due to pressure imbalances (between the internal epoxy and the cooling water) deforming the shell.

#### B. Short Fibre Infused

- 1) Advantages:
- 1) Same as Section III-A.1 1 (Quick print).
- 2) Same as Section III-A.1 2 (Save material).
- 3) Same as Section III-A.1 3 (Strong print).
- 4) The reinforcement method allows for a stronger blade due to the added glass fibre.
- 5) The reinforcement is multidirectional.
- 6) Same as Section III-A.1 5 (Removable shell).
- 2) Disadvantages:
- 1) Same as Section III-A.2 1 (Warping).
- 2) Same as Section III-A.2 2 (Exo-therm).
- 3) Fibre displacement during infusion can occur due to the unconstrained fibres shifting from the flow of the infused resin.
- 4) Since the short fibres are loosely stacked inside the printed shell, with no precise fibre placement, there is a lack of control of fibre direction.
- 5) Similar to Item 4 there is also a lack of control of fibre density.
- 6) Similar to Item 4 there is also a lack of control of fibre distribution.
- Similar to Section III-A.2 5 as there is a lack of precise control of blade volume during the infusion process (Unbalancing).

#### C. Pultruded Rod Reinforced

- 1) Advantages:
- 1) Good geometry accuracy (no shell warping) due to the internal support structure.

- 2) The blade is relatively light weight since the majority of the blade is the porous printed structure.
- 3) Little resin is required in this process which saves on costs.
- 4) Very strong blade in primary loading direction (along blade length) due to directional fibres of the pultruded rods.
- 5) No exo-therm runaway problems since little resin is used in this process.
- 2) Disadvantages:
- Complicated printing since the pultruded rods must run inside the printed matrix from blade root to tip. This also limits any curving of the blade as well as limiting the minimum volume of the blade (to be able to accommodate the size of the pultruded rods).
- 2) This process requires long print times (12-24 hours) due the extra complexity and material of the internal support structure.
- 3) Related to Item 2, extra print material is required for this process which also increases costs.
- Delicate shell structure since there is no reinforcement of profile other than the printed matrix and some pultruded rods.
- 5) Unbalancing between blades weights if the epoxy casting of the blade root and tip are not carefully controlled.

#### **IV. CONCLUSIONS**

This general method of manufacturing small wind turbine blades using FDM assisted composite structures has proven to be very successful. The process has a short turn around time capable, from new design to final product, of manufacturing a set of 3 blades within 5-7 days. The process also has proven to be cost effective, with the cost per blade varying from about R150 to R500 depending on size and reinforcement method. The accuracy of the blade geometry has also proven to be sufficient for blade manufacturing. Finally, the practical testing of the blades was completed at Kestrel's outdoor wind tunnel with wind speeds of up to 9 m/s (32 km/h) average and turbine rotational speeds of over 500 rpm (Figure 14). Of the three compared manufacturing methods (Section III), the most favourable was the pultruded rod reinforced method as this produced the most geometrically accurate blade which was relatively light weight. Further testing of the blades is under way with an in-field mounted setup capable of long term testing (for energy yield and fatigue results). Further research is also recommended on 3D printing moulds for composite layup manufacturing of small turbine blades.

#### REFERENCES

- Sean N Poole. Optimisation of a Mini Horizontal Axis Wind Turbine to Increase Energy Yield During Short Duration Wind Variations. Unpublished as of 2015.
- [2] Pramod Jain. Wind Energy Engineering. McGraw-Hill, New York, 2011. ISBN 9780071714778.
- [3] Repetier. Repetier. URL http://www.repetier.com/. Date accessed: 21 Oct 2015.



Fig. 14: Testing of Manufactured Wind Turbine at Kestrel's Outdoor Wind Tunnel

- [4] Tony Burton, David Sharpe, Nick Jenkins, and Ervin Bossanyi. Wind Energy Hand-book. John Wiley & Sons, Ltd, 2002. ISBN 9780470846063.
- [5] Scott Davidson. Grasshopper 3d. URL http://www.grasshopper3d.com/. Date accessed: 21 Oct 2015 Robert McNeel & Associates. Rhino
- [6] Robert Robert McNeel & Associates. Rhino url:https://www.rhino3d.com/, date accessed: 21 Oct 2015 3D.
- [7] D.I.Y Electronics. DIYElectronics Prusa I3 Build Guide 2014. Technical report, 2014.
- [8] Bart Meijer. Marlin User Guide. Technical Report March, 2014.

# Gait Adaptation of a Six Legged Walker to Enable Gripping

Tracy Booysen<sup>1</sup> and Florian Reiner<sup>2</sup>

Abstract— The value of researching legged walkers can be seen as they are being used more often. Walking robots are better at navigating over rough terrain and, depending on the leg configuration, can have redundancy to mitigate leg failure. As a four legged walker is statically stable, a six legged walking robot has are two redundant legs. This addition of extra legs allows for either a six legged walking gait or a four legged walking gait where two of the legs are utilised as grippers. In this paper various gripper types are analysed and the affect of the performance with regards to walking stability and gait design are analysed.

#### I. INTRODUCTION

Legged robots have significant advantages over tracked or wheeled robots in that they have more stability in cluttered or rough terrain [1]. They can also provide redundancy by increasing the number of legs on the platform. To achieve static stability at least three legs must be on the ground at any given time [1], [2]. This means that any robot with a leg count greater then four can be statically stable. Increasing the number of legs increases the walking speed and decreases the portion of mass that each leg has to carry. However each additional leg adds to the computational and control complexity of the robot. A six legged walking robot allows for some leg redundancy at an increased walking speed while keeping the cost and complexity relatively low.

In any robotic field, the usability of the robot depends on the actuators and sensors attached to it. For a pure inspection robot, cameras may be sufficient to achieve a variety of operations. However in many cases the ability to grasp or manipulate an object greatly increases the usefulness of a robot. This has led to a number of robots being equipped with dedicated arms and grippers. However the addition of a dedicated gripper adds additional mass, power consumption and control complexity to the system. Depending on the gripper system required, the redundant legs on a six legged walker could be transformed into grippers while maintaining static stability.

This motivates the ongoing project presented in this paper. This paper describes the modification of an existing six legged walker to include two leg/gripper combinations. This allows the robot to approach an object using the faster six legged gait, transition into a four legged gait, grip an object and walk whilst still holding the object. Various gripping surfaces are compared and gripping performance

<sup>2</sup>Florian Reiner is with the Department of Mechanical Engineering, University of Cape Town, Cape Town, South Africa RNRFL0001@myuct.ac.za

using different size and shaped objects is analysed. To allow the robot to be truly adaptive, a new four legged gait was developed and the effect of the new gait on the platform stability is investigated.

The paper begins by describing the current hexapod before the design and implementation of the gripper is presented. Then the gripper control and gait modification techniques are introduced. A comparison of some performance is presented and we conclude with future work in Section VII.

#### II. CURRENTLY DESIGNED HEXAPOD

In 2011 the University of Cape Town developed a six legged walker with three degrees of freedom per leg [3]. Each leg contained three RX-28 dynamixel motors which allowed for monitoring of torque, temperature and various other parameters. In its circular configuration the robot is 850 mm in diameter and weighs 4.4 kg. The original robot can be seen in Figure 1(a).

The original robot was designed to be reconfigurable and thus to have two possible shapes. The six legs could be arranged around either a circular or a rectangular body. The rectangular body made the walking gait simpler as each leg swung in the plane of forward motion. However the advantage of the circular body was the complete omnidirectionality of the system as the robot has no front, back, left or right. The control of the robot included a directional change ability so that an operator could redefine any side of the robot as the front and continue moving forward in the newly defined direction. Another advantage noted by Takahashi et al [4] of arranging legs equally around a circular body was to reduce the interference between legs and maximize working space for any manipulators on the robot. Because of this most of the development focussed on implementation on the circular body in the assumption that reworking any control into the rectangular body would be a fairly simplistic task.

#### A. Similar Robots

Between 1995 and 2002 a similar hexapod was developed by a Japanese research team led by F. Nagoya. This hexapod used an "integrated limb mechanism" to adapt two legs for use as grippers while stationary. Two prototypes were designed, Melmantis-1 and Melmantis-2 which demonstrated the basic gripping of objects. However the robots could only grip while stationary and the physical design of the robot was limited as they were still in proof of concept phases [5], [6], [7].

Another six legged walking robot that has an integrated gripper is the LAURON which was designed as a walking robot for the planetary exploration spacebot cup competition

<sup>&</sup>lt;sup>1</sup>Tracy Booysen is with the Department of Mechanical Engineering, University of Cape Town, Cape Town, South Africa tracy.booysen@uct.ac.za

[8]. It was a biologically inspired robot and each leg has four degrees of freedom. By moving the weight onto the rear legs and equipping both front legs with grippers, LAURON is capable of using both front legs for gripping tasks. Each leg can grip up to 2.3 kg and up to 50 mm in diameter. LAURON has a storage area that is used for store gripped objects whilst the robot is walking.

While walking robots with grippers are not uncommon, very few robots are able to walk while gripping an object. One of these robots, developed by some of the same researchers as Melmantis, was named ASTERISK [4] which used a circular six legged walker with every limb being able to be a gripper. Limbs directly opposite to each other were used for gripping and to ensure stable walking in every orientation, the legs doing the gripping had to be changed. This robot was further adapted to use adjoining limbs and to modify its standing position when in gripping mode. The robot then places the object onto its body and continues on in a six legged mode [9].

#### **III. GRIPPER DESIGN**

Two types of grippers were initially considered for use; frictional and encompassing. Frictional grippers rely solely on the frictional force generated from the clamping force between two plates. The amount of frictional force generated is proportional to the clamping force and the frictional coefficient of the clamping material. Encompassing grippers cradle the object so that it rests on part of the gripper. Because it is not reliant on friction, minimal to no clamping force is required to grip the object [10]. Although encompassing grippers use less power due to the low clamping force and objects with a very low co-efficient of friction can be lifted, the shape of the gripper must match the shape of the object. To increase the number of unique objects that the robot could lift, a frictional gripper was designed.

To maximise the area of gripping and to increase the gripping force, a gripper must remain normal to the object it is gripping. Due to the circular nature of the robot the gripper needed to have two degrees of freedom. This was because the angle of each side of the gripper changes as the legs close. The addition of any extra motors or actuators to provide these additional degrees of freedom added too much to the control and cost overhead. A simpler solution was to design the gripper with a rod-end joint. This joint acts similarly to a universal joint and allows extra degrees of freedom. It is also stiffer then a universal joint and will remain in position after being moved.

To create the gripper a thin layer of rubber was glued to a thicker layer of foam. The rubber provided a higher coefficient of freedom then the foam and the foam encompassed some of the object. This was attached to the robot with Velcro so that different gripping pads using different material could be investigated. Two of these grippers were attached to two adjacent legs of the hexapod.

A render of the gripper can be seen in Figure 1(b).

#### A. Gripper Control

Other than the position of the gripper relative to the object, the most important variable to control is the gripping force. If the force is too low, then the object will slip out of the gripper. If the force is too high, then there is a risk of damaging either the motors themselves, or the object being gripped.

Because each leg was equipped with Dynamixel motors, torque can be continuously monitored. The clamping force is directly related to the torque on the horizontal motors of the grippers. When in gripping mode, software continually monitors these motors and compares the current torque value with a safety threshold. If the current torque exceeds this threshold the horizontal position of the motors will be held constant and the grippers will not close any further. The safety threshold is defined by the user and this allows the operator to preselect an appropriate clamping force depending on the object being lifted.

In addition the vertical gripper motor torques and the torques on the front legs are also continuously monitored as lifting heavy objects will increase the torques applied to these motors as well. If the maximum torque is reached the robot will not continue to lift the mass. Figure 1(c) shows how the gripper would grip an object such as a football. Additionally the position of the front and rear legs can be seen. In order to stabilise the platform during gripping the robot had to shift the leg positions to mimic a four legged walker.

#### B. Gripper Performance

With the motors powered at their maximum voltage of 14.4 V the performance of the gripper was analysed, the results of which can be seen in Table I. The minimum object size was limited by the size of the gripper pads. If an object less than half the size of the gripper pads was lifted, the gripper would tend to buckle inwards while lifting the object. This did not always occur, but objects under 30 mm were not lifted reliably. Should smaller objects need to be lifted the size of the gripper pad should be reduced accordingly. Many different sizes and shaped of objects were able to be lifted by the grippers, some of which can be seen in Figure 1(d). 230 mm was the largest size object that could be lifted and this was limited by how wide the front legs could open. This value was encoded as a hard limit into the software to avoid collisions with other legs of the robot.

The grippers were able to lift a variety of surface materials including wood, rough metal, polished metal, rubber, foam and plastics. There was no noticeable difference between materials of a similar mass. The high coefficient of friction provided by the rubber pads was a larger factor than the coefficient of friction provided by the object.

Several non-parallel items and items without flat surfaces were gripped. While items with slanting sides or distinct corners were lifted successfully, non-uniform objects like stones were not lifted reliably.

The gripper was also evaluated with three different gripping surfaces, rubber only and two different densities of



(a) Complete Robot



(b) The new gripper



(c) Robot Gripping a Ball



(d) All the objects sucessfully gripped

Fig. 1. The Robot and Gripper Operation

TABLE I Gripper Geometry Results

	Min Diameter	Max Diameter	Max Mass at
	/Length (mm)	/Length (mm)	full extension
Rectangular	30 mm	230 mm	780 g
Prisms			-
Cylinders	30 mm	>150 mm	600 g
Spheres	30 mm	220 mm	>400 g

foam. The maximum mass that the gripper could lift at minimum, mid and maximum extension was measured and is shown in Table II.

Unsurprisingly, far greater masses could be lifted if the gripper arms were closer to the body, this is because the raising torque directly increases with increasing distance due to the longer moment arm. The pads with only rubber proved to be the most successful; mainly because the foam pads would tend to buckle eventually at higher masses.

The results also showed that at all three distances the high density foam performed better than the lower density foam, which was presumably due to the higher stiffness being more resistant to buckling.

At masses of more than 1.5 kg, the front legs of the hexapod would tend to overheat rapidly and become overloaded. In the gripping position, the front legs carry the majority of the weight of the robot in addition to the gripped weight. Therefore to prevent overloading of the motors whilst still testing the capabilities of the gripper, the robot was rested on a stand for the tests at minimum extension.

TABLE II GRIPPER MASS RESULTS

	Min	Mid	Max
	Extension	Extension	Extension
	(60 mm)	(140 mm)	(210 mm)
Rubber pads	2070 g	1210 g mm	780 g
without foam			
Rubber pads and	1720 g mm	1050 mm	670 g
SPX 25 foam			
Rubber pads and	1790 g mm	1080 mm	710 g
SPX 60 foam			

#### IV. GAIT MODIFICATION AND STABILITY TRIANGLES

Because two of the six legs are grippers, when the robot entered gripping mode only four of the legs remained for walking. This meant that a new gait would have to be developed. The standard four legged walker gait would entail keeping three legs on the ground to create a stable triangle and then moving one leg at a time. However in this case it was more complex due to both the geometry of the circular robot and the mass of the object being gripped. Once the front two legs had gripped an object, the robot's Centre of Gravity (CG) moved significantly forward. This caused the robot to pitch over. To prevent this from happening the front two legs had to move further forward and the rear legs further backwards.

While this worked while the robot was stationary, once the first leg was lifted, the CG moved further forwards and the robot began to pitch. This is because the four legs are not arranged equally around the CG. Although this problem was exacerbated by the circular body, it should still appear in a rectangular hexapod. In addition, the larger the mass that is lifted, the larger this problem becomes.

Two solutions were implemented to stabilise the platform. Firstly, in order to allow for maximum stability, the legs must have at least a  $180^{\circ}$  range of motion ( $90^{\circ}$  from the centre to the front of the robot and  $-90^{\circ}$  from the centre to the rear of the robot). This allows the front legs to move as far forward as is possible and the rear legs to move as far back as possible. The circular body was machined to increase the range of motion which in turn provided additional stability. However this still did not solve the pitching problem.

In order to be able to lift legs without the robot pitching over, the CG had to be altered by sliding the body of the robot forward. This was done by moving the body of the hexapod forward while leaving the legs in the same position. Figure 2(a) shows the position of the CG (shown in white) and the position of the body in relationship to the legs before a slide was implemented. Figure 2(b) shows the same items, but after the slide was implemented. As can be seen this effectively moves the CG from the rear (where it would be safe to lift a front leg) to the front (where it would be safe to lift a rear leg).



Fig. 2. Positioning of CG Before and After the Sliding Motion

The alteration of the gait caused the stability of the robot to change when compared with the six legged gait. When the hexapod walked in six legged mode, the pitch of the robot showed the greatest deviation, changing by a total of  $3^{\circ}$  [3]. When the hexapod walked in four legged walking mode the pitch deviated by  $8^{\circ}$ . This is shown in Figure 3. The peaks of this instability occur when the rear leg of the hexapod is lifted. This also coincides with the peaks of the roll - which was approximately  $4^{\circ}$ . This effect could be lessened by extending the gripper arms outwards, however at that point the lifting of the front leg would cause a pitching moment. The torque on the front legs (when not gripping an object) was raised from 15.77 kg/cm in six legged walking mode to 18 kg/cm in four legged walking mode.



Fig. 3. Pitch of Hexapod While in Four Legged Mode

#### V. AUTOMATIC PITCH DETECTION WHILE WALKING

Altering the gait and body allowed the robot to walk whilst holding a specific object. However there was a limitation on the mass of the object that can be held and the position it can be held in whilst walking. These limitations are not the same as the masses that the gripper can lift, as lifting objects is done with four legs on the ground and walking with three. If too large a mass was lifted and the grippers extended too far out, the robot would pitch over when walking. Since the gait was not an adaptive one, there needed to be some form of pitch detection with some form of reaction to maintain stability.

For this application stability was defined as the ability of the robot to remain upright under all circumstances. For the robot to remain upright, the body of the robot must remain parallel to and not touch the ground. If any part of the robot's body has touched the ground then it will be considered to be unstable. Because this is to be avoided at all costs some method of detecting unstable positions had to be included in the system. This would be active during both the gripping motion and the walking motion.

When in gripping mode (walking or stationary) the only way for the robot to tip would be for it to pitch forwards or backwards. This indicated that the pitch of the robot should be affected by instability much more then roll or yaw. An IMU was placed on the hexapod and the gripper was extended outwards until the robot pitched over. The pitch of the robot was recorded and can be seen in Figure 4. Note that in tests 1-5 the front left leg was off the ground and in test 6-10 the front right leg was off the ground. A single leg was raised to further mimic the position of the legs whilst walking.

**Pitch during Falling Tests** 



Fig. 4. Pitch of Hexapod as Tipping Occurs

These results show that tipping is easily detected when the IMU reads a value of  $-4^{\circ}$ . However during walking the average pitch of the robot is greater so a value in the range of  $-8^{\circ}$  was chosen.

In order to prevent a loss of stability, the robot was required to react to the detection of tipping and prevent it from continuing. Three possible reaction schemes were considered:

- Leaning Back: If the hexapod is tipping forwards then this could be counteracted by leaning the body backwards. This would be achieved by raising the front legs and lowering the rear legs, thus shifting the CG backwards and stabilizing the system.
- Retracting the Grippers: If the object is too heavy then the hexapod will tend to tip forwards. Retracting the gripper will minimise the moment arm and bring the CG back towards the centre of the hexapod.
- Dropping the Item: Dropping the item (if fast enough) will prevent the hexapod from tipping.

Although leaning back would seem to be the most natural response to implement, the hexapod did not tip over while gripping an object. The limiting factor was the fact that the two front-most legs vertical motors became overloaded before tipping occurred. Tipping only occurred during walking and this response would not be possible to implement whilst in-gait.

During the four legged gait if the grippers were extended, even without holding an object, the robot would tip. This is due to the geometry of the robot being very sensitive to changes in the CG. Extended arms (with or without a mass) shift the CG too far forward to allow walking. Thus retracting the grippers towards the body would not solve the tipping problem as they would always be retracted when walking.

The final solution implemented was to drop the object if the robot became unstable whilst walking. This would only occur if the object was particularly heavy. Because the mass of the object could be inferred from the vertical gripper motor torques the IMU is removed once the maximum stable mass to be lifted was found.

#### VI. CONCLUSION

Overall the addition of the grippers increased the functionality of the robot. The six legged walking gait was not affected by the addition of the gripping pads and objects of all shaped and sizes were able to be lifted. Although a four legged walking gait was successfully implemented, the increase in the pitch and roll of the body increased to a more undesirable level.

The gripping force was able to be adjusted to prevent damage to the motors and the monitoring of the torque allowed for automatic pitch prevention to be implemented.

The additional mass added to the system when objects were being gripped added and increased torque to the vertical motors. This caused overheating when lifting heavy objects or if the robot paused during a gait with only a few legs on the ground. The motors came equipped with a thermal shut down, however this often failed to work and caused motors to become permanently damaged.

#### VII. FUTURE WORK

The performance of the hexapod in its rectangular configuration should be evaluated in order to see whether the shifting CG problem is lessened. An analysis on the functionality of circular versus rectangular robots should be considered however, because circular walkers have greater freedom of motion for gripping and have no defined front. Therefore care should be taken before moving to the simpler, rectangular shape.

Because the implementation of the four legged walk increased the pitch of the body, additional gait analysis and modification should be performed. An ideal gait would either maintain perfect stability of the body to enable accurate camera vision and advanced functions like path planning or would sacrifice stability for speed of motion.

An additional degree of freedom could be added to the grippers which would allow the mass of the object to be distributed to the centre of the hexapod. This would solve much of the shifting CG problem. Objects could also be dropped off and retrieved later.

Stronger motors are needed on the vertical joints of all the legs. This would allow for larger objects to be carried while walking and while stationary. While the motors rarely became overloaded and were protected from harm by implementing a thermal shut down procedure, failure of the vertical motors still occurred. This could be mitigated by using higher torque motors.

A simulation in an environment such as Gazebo would be ideal as the robot has the potential to damage itself when walking if the gait is implemented incorrectly.

Currently larger motors are being added to the vertical joints of each leg. In addition force sensors are being added to the feet to determine when they are on the ground. Once this has been completed further work in closing the control loop using the foot sensors could be completed. In addition, further processing power is being added to make each leg a modular, stand-alone, closed-loop system, communicating with the main controller as a unit, rather as 18 individual motors.

#### REFERENCES

- D. C. Kar, "Design of a statically stable walking robot: A review," Journal of Robotic Systems, pp. 671–686, 2003.
   G. Carbone and M. Ceccarelli, "Legged robotic systems," in Cutting
- [2] G. Carbone and M. Ceccarelli, "Legged robotic systems," in *Cutting Edge Robotics*. Pro Literatur Verlag, 2005.
- [3] T. Booysen and S. Marais, "The development of a remote controlled, omnidirectional six legged walker with feedback," in *IEEE Africon*, 2013.
- [4] Y. Takahashi, T. Arai, Y. Mae, K. Inoue, and N. Koyachi, "Development of multi-limb robot with omnidirectional manipulability and mobility," in *Proceedings of the 2000 IEEE International Conference* on Intelligent Robotics and Systems, (IROS 2000). IEEE, 2000.
- [5] N. Koyachi, T. Arai, H. Adachi, and M. A. Adachi, "Design and control of hexapod with integrated limb mechanism: Melmantis," in *Proceedings of the 1996 IEEE International Conference on Intelligent Robotics and Systems, (IROS 1996).* IEEE, 1996, pp. 877–882.
- [6] N. Koyachi, T. Arai, H. Adachi, A. Murakami, and K. Kawai, "Mechanical design of hexapods with integrated limb mechanism: Melmantis-1 and melmantis-2," in *Proceedings of the 1997 IEEE International Conference on Advanced Robotics, (ICAR 2000).* IEEE, 1997, pp. 279–278.
- [7] N. Koyachi, H. Adachi, M. Izumi, and T.Hirose, "Control of walk and manipulation by a hexapos with integrated limb mechanism: Melmantis-1," in *Proceedings of the 2002 IEEE International Conference on Robotics and Automation, (ICRA 2002).* IEEE, 2002, pp. 3353–3358.
- [8] G. Heppner, A. Roennau, J. Oberlander, S. Klemm, and R. Dillmann, "Laurope - six legged walking robot for planetary exploration participating in the spacebot cup," in *Proceedings of the 2015 International Conference on Automation and Robotics in Space*, (ASTRA 2015), 2015.
- [9] T. Takubo, T. Arai, K. Inoue, H. Ochi, T. Konishi, T. Tsurutani, and Y. Hayashibara, "Integrated limb mechanism robot asterisk," *Journal* of Robotics and Mechatronics, pp. 203–214, 2006.

[10] "Robotic gripper sizing: The science, technology and lore," http://www.grippers.com/size.htm, last visited on April 2014.

## Semi-supervised Spectral Connectivity Projection Pursuit

David Hofmeyr<sup>1</sup> and Nicos Pavlidis<sup>2</sup>

Abstract—We propose a projection pursuit method based on semi-supervised spectral connectivity. The projection index is given by the second eigenvalue of the graph Laplacian of the projected data. An incomplete label set is used to modify pairwise similarities between data in such a way that penalises projections which do not admit a separation of the classes (within the training data). We show that the global optimum of the proposed problem converges to the Transductive Support Vector Machine solution, as the scaling parameter is reduced to zero. We evaluate the performance of the proposed method on benchmark data sets.

#### I. INTRODUCTION

Projection pursuit is a data driven optimisation problem, defined as follows. For a data set  $X = \{x_1, ..., x_N\}$  in  $\mathbb{R}^d$ , optimise over the set of unit-norm vectors,  $\{v \in \mathbb{R}^d | ||v|| = 1\}$ , a predefined measure of quality of the projected data set,  $v \cdot X = \{v \cdot x_1, ..., v \cdot x_N\}$ . These unit-norm vectors are referred to as *projection vectors*, or simply *projections*, and the measured quality of the projected data set is referred to as the *projection index*.

Semi-supervised classification refers to the construction of a classifier, i.e., a map from the data space to a set of class labels, using a set of "training" data whose true class labels are known as well as a set of "test" data whose labels are to be inferred from the classifier. In supervised classification, on the other hand, only the training data and associated labels are used in the construction of the classifier. In using a classifier for class prediction there is an implicit assumption that the distribution of the test data resemble somewhat that of the training data, and therefore utilising spatial distribution information of the test data might be useful in better predicting their class memberships.

There are numerous approaches to the problem of semisupervised classification, see [1] for a recent review of standard methods. Underlying many of these methods is the so-called *cluster assumption*; that different classes manifest single clusters, and so can be separated by data sparse regions. Of these methods, arguably the most popular are those based on Transductive Support Vector Machines (TSVMs). The original TSVM problem [2] is formulated as follows. Given labelled data  $\mathbf{X}^L = \{x_1, ..., x_l\}$  with labels  $Y^L \in \{-1, +1\}^l$  and unlabelled data  $\mathbf{X}^U = \{x_{l+1}, ..., x_{l+u}\}$ , find  $Y^U \in \{-1, +1\}^u$  s.t. a Support Vector Machine (SVM)

and Statistics, Lancaster University, Lancaster, UK, LA1 4YF. d.hofmeyr@lancaster.ac.uk

 $^2\rm Nicos$  Pavlidis is with the Department of Management Science, Lancaster University, Lancaster, UK, LA1 4YX.

classifier trained on  $\mathbf{X}^L \cup \mathbf{X}^U, Y^L \cup Y^U$  achieves the largest margin. This is a combinatorial problem and difficult to solve for any reasonably sized data set. Approximations based on local search heuristics [4] or continuous relaxations are solved instead [3], but these are highly susceptible to local optima.

Accepting the cluster assumption leads us to consider semi-supervised methodology that is consistent with popular notions of clusterability. Spectral clustering has become increasingly popular due to its strong performance in a variety of application areas [5]. In spectral clustering, clusters are defined as strongly connected components of a graph defined over the data in which edge weights assume values equal to the similarity between the adjacent vertices. A continuous relaxation of the minimum ratio cut problem is solved, using the eigenvectors of the graph Laplacian matrix. Recently a projection pursuit method for learning the projection along which a set of data are minimally connected under this cluster definition was proposed [6]. The authors show that the projection along which the data are minimally connected converges to the vector normal to the largest margin hyperplane through the data, as the scaling parameter is reduced to zero. In this paper we extend this work to include partial supervision via an incomplete set of labels, as in semi-supervised classification. We show that if the labels are incorporated in a specific way, then the convergence result of [6] extends to the semi-supervised setting, i.e., the optimal projection for semi-supervised spectral connectivity converges to the vector normal to the optimal TSVM hyperplane. This establishes an asymptotic connection between our proposed method and popular semi-supervised classification methods.

The Remainder of this paper is organised as follows. In Section II we give a brief introduction to spectral clustering. In Section III we introduce the proposed methodology. We provide experimental results in Section IV and give some concluding remarks in Section V.

#### II. SPECTRAL CLUSTERING

In this section we give a very brief introduction to spectral clustering, with particular attention to binary partitioning. For a thorough introduction, the reader is directed to [5]. Let  $X = \{x_1, ..., x_N\}$  be a given data set. The ratio cut problem is defined as follows,

$$\min_{C \subset X} \sum_{\substack{i,j:x_i \in C\\x_j \notin C}} \text{similarity}(x_i, x_j) \left( \frac{1}{|C|} + \frac{1}{|X \setminus C|} \right), \quad (1)$$

David Hofmeyr's Ph.D is funded by the Engineering and Physical Sciences Research Council (EPSRC) and the Oppenheimer Memorial Trust. <sup>1</sup>David Hofmeyr is with the Department of Mathematics

where  $|\cdot|$  is the cardinality operator. The similarity between two points is generally determined by a non-negative, decreasing function,  $k : \mathbb{R} \to \mathbb{R}^+$ , of the distance between them. The above problem can be formulated in terms of the *Laplacian matrix*, L := D - A, where A is the *affinity matrix*, with  $A_{ij} = \text{similarity}(x_i, x_j)$ , and the diagonal matrix D is called the *degree matrix*, with  $D_{ii} = \sum_{j=1}^{N} A_{ij}$ . For  $C \subset X$ define the vector  $f^C \in \mathbb{R}^N$  such that,

$$f_i^C = \begin{cases} \sqrt{|X \setminus C|/|C|}, & x_i \in C\\ -\sqrt{|C|/|X \setminus C|}, & x \notin C. \end{cases}$$
(2)

Then (1) can be written as,

$$\min_{C \subset X} f^C \cdot L f^C \text{ s.t. } f^C \perp \mathbf{1}, \|f^C\| = \sqrt{N}.$$
(3)

The above problem is NP-hard [7], and so instead a continuous relaxation, in which the discreteness condition on the vector  $f^{C}$  (2) is relaxed, is solved instead. The solution to the relaxed problem is given by the second eigenvector of L. The second eigenvalue therefore provides a lower bound for the normalised aggregated similarities from pairs of data belonging to different elements of the optimal partition, arising from the solution to (1). While the optimal solution to the relaxed problem can induce partitions which are arbitrarily far from the optimal solution to (3) [8], in many practical applications the solutions tend to be similar. Furthermore, in the univariate case the resulting partition tends to be very similar to the true optimum. Obtaining the projection which minimises the second eigenvalue of the Laplacian therefore tends to result in projections along which the optimal partition arising from (1) is loosely connected, i.e., the elements of the partition are well separated.

#### III. METHODOLOGY

In this section we provide details for how to find locally optimal projections for bi-partitioning based on semisupervised spectral clustering. We formulate the problem as a projection pursuit, where the projection index is given by the second eigenvalue of the Laplacian of the projected data. We assume we have a set of N data, l of which have labels in  $\{-1, +1\}$  which define their class membership, and u = N - l are unlabelled.

Following the method of [6] we formulate the projection vectors in terms of their polar coordinates. Let  $\Theta = [0, \pi)^{d-2} \times [0, 2\pi)$  and for  $\boldsymbol{\theta} \in \Theta$ , define the projection vector  $v(\boldsymbol{\theta})$  by,

$$v(\boldsymbol{\theta})_i = \begin{cases} \cos(\boldsymbol{\theta}_i) \prod_{j=1}^{i-1} \sin(\boldsymbol{\theta}_j), & i = 1, \dots, d-1 \\ \prod_{j=1}^{d-1} \sin(\boldsymbol{\theta}_j), & i = d. \end{cases}$$
(4)

We will use the following notation. For  $\boldsymbol{\theta} \in \Theta$ , we write  $L(\boldsymbol{\theta})$  for the Laplacian of the projected data set  $P(\boldsymbol{\theta}) := v(\boldsymbol{\theta}) \cdot X$ , and  $\lambda_2(\boldsymbol{\theta})$  for the second eigenvalue of  $L(\boldsymbol{\theta})$ . If the similarities between pairs of projected data are Lipschitz and continuously differentiable functions of  $\boldsymbol{\theta}$ , then  $\lambda_2(\boldsymbol{\theta})$  is Lipschitz and continuously differentiable almost everywhere [6]. This allows us to find local optima via generalised gradient descent. The derivative of  $\lambda_2(\boldsymbol{\theta})$  with respect to  $\boldsymbol{\theta}$ 

can be decomposed using the chain rule into the product  $D_{P(\theta)}\lambda_2(\theta)D_{v(\theta)}P(\theta)D_{\theta}v(\theta)$ , where  $D_{\cdot}$  is the differential operator. Derivations of the following can be found in [6]. If  $\lambda_2(\theta)$  is a simple eigenvalue, then

$$\frac{\partial \lambda_2(\boldsymbol{\theta})}{\partial P(\boldsymbol{\theta})_k} = \frac{1}{2} \sum_{ij} (u_i - u_j)^2 \frac{\partial A(\boldsymbol{\theta})_{ij}}{\partial P(\boldsymbol{\theta})_k},\tag{5}$$

where  $A(\boldsymbol{\theta})$  is the affinity matrix of the projected data set. The matrix  $D_{v(\boldsymbol{\theta})}P(\boldsymbol{\theta}) \in \mathbb{R}^{N \times d}$  has *i*-th row  $x_i^{\top}$ , and the matrix  $D_{\boldsymbol{\theta}}v(\boldsymbol{\theta}) \in \mathbb{R}^{d \times (d-1)}$  has *i*, *j*-th element,

$$\frac{\partial v(\boldsymbol{\theta})_i}{\partial \boldsymbol{\theta}_j} = \begin{cases} 0, & i < j \\ -\sin(\boldsymbol{\theta}_j) \prod_{k=1}^{j-1} \sin(\boldsymbol{\theta}_k), & i = j < d \\ \cos(\boldsymbol{\theta}_j) \cos(\boldsymbol{\theta}_i) \prod_{k < i, k \neq j} \sin(\boldsymbol{\theta}_k), & j < i < d \\ \cos(\boldsymbol{\theta}_j) \prod_{k \neq j} \sin(\boldsymbol{\theta}_k), & i = d. \end{cases}$$
(6)

What remains is to address the similarity function. Within spectral clustering, pairwise similarities between data are defined by a decreasing function of the distance between them. That is, similarity $(x_i, x_j) = k(d(x_i, x_j))$ , where k : $\mathbb{R}^+ \to \mathbb{R}^+$  is decreasing and  $d(\cdot, \cdot)$  is a metric. The function k often takes the form of a kernel, and we use the Gaussian kernel, given by

$$k(x) = \exp\left(-\frac{x^2}{2}\right).$$
(7)

Controlling the balance of the partition, i.e., the relative sizes of the resulting clusters, is an important feature in semisupervised classification [1]. We control this balance in the same way as in [6] within the metric  $d(\cdot, \cdot)$ . In particular, for a univariate data set P,

$$d(P_i, P_j) := \frac{|T(P_i) - T(P_j)|}{\sigma}, \tag{8}$$

where  $\sigma > 0$  is the *scaling parameter*, and the function T is used to decrease the distance between points lying outside a chosen interval [m, M] and other points, to induce more balanced splits.

$$T(x) := \begin{cases} -\delta \left( m - x + (\delta(1-\delta))^{\frac{1}{\delta}} \right)^{1-\delta} \\ + \delta(\delta(1-\delta))^{\frac{1-\delta}{\delta}}, & x < m \\ x - m, & x \in [m, M] \\ \delta \left( x - M + (\delta(1-\delta))^{\frac{1}{\delta}} \right)^{1-\delta} \\ - \delta(\delta(1-\delta))^{\frac{1-\delta}{\delta}}, & x > M \end{cases}$$
(9)

We set  $m = \mu_P - \beta \sigma_P$  and  $M = \mu_P + \beta \sigma_P$ , where  $\mu_P$  and  $\sigma_P$  are the mean and standard deviation of P respectively and  $\beta$  is used to control the width of the interval [m, M]. See [6] for details on the effect of the function T. The parameter  $\delta$  takes values in (0, 0.5], with smaller values increasing the similarity of points outside [m, M] with other points to a greater degree. The value of this parameter does not play a huge role in performance [6].

So far we have not discussed how we incorporate label information into this framework. The ratio cut is

inherently a connectivity based partitioning method, and though the spectral clustering solution is a relaxation, its behaviour mimics this connectivity property. We wish to use these labels to modify the pairwise similarities in such a way that projections which do not admit a separation of the (known) classes are penalised. By this we mean projections for which  $\exists i, j$  s.t.  $y_i = +1, y_j = -1$  but  $P(\boldsymbol{\theta})_i < P(\boldsymbol{\theta})_j$ , i.e., the positive labelled projections do not all lie above all negative labelled projections. This can be achieved by ensuring that along any such projection there is a chain of high pairwise similarities connecting the entire projected data set, as follows. For brevity we temporarily drop the notational dependence on  $\theta$ . Define,

$$A_{ij} = \begin{cases} k(d(P_i, P_j)), & x_i, x_j \in \mathbf{X}^U \\ k(d(P_i, P_j)) + ((P_i - P_j)^+)^{1+\epsilon}, & y_i = -, y_j \neq - \\ k(d(P_i, P_j)) + ((P_j - P_i)^+)^{1+\epsilon}, & y_i = +, y_j \neq + \\ H, & y_i = y_j, \end{cases}$$
(10)

where  $(x)^+ = \max\{0, x\}$  and  $H \ge 1$  is a chosen constant which affects the influence of the known labels. We use  $y_i \neq j_i$ +1 (resp.  $y_i \neq -1$ ) to mean that  $y_i = -1$  or  $x_i$  is unlabelled (resp.  $y_i = +1$  or  $x_i$  is unlabelled). The exponent  $1 + \epsilon$ , where  $\epsilon$  is some small positive number, ensures continuous differentiability of the associated additions while having a practical influence much like the hinge loss function. We'll refer to these additions as penalties. With the above formulation we can derive expressions for  $\partial A_{ij}/\partial P_k$  for all i, j, k. If  $x_i, x_j \in \mathbf{X}^U$  then,

$$\frac{\partial A_{ij}}{\partial P_k} = \frac{\partial A_{ij}}{\partial T(P_j)} \frac{\partial T(P_j)}{\partial P_k} = \frac{\partial k(d(P_i, P_j))}{\partial T(P_j)} \frac{\partial T(P_j)}{\partial P_k}$$

$$= \frac{T(P_i) - T(P_j)}{\sigma^2} \exp\left(-\frac{d(P_i, P_j)^2}{2}\right) \frac{\partial T(P_j)}{\partial P_k}.$$

If  $y_i = +1, y_j \neq +1$ , then if  $P_i > P_j$  we have the same formulation as above. Otherwise,

$$\frac{\partial A_{ij}}{\partial P_k} = \frac{\partial k(d(P_i, P_j))}{\partial T(P_j)} \frac{\partial T(P_j)}{\partial P_k} + \frac{\partial (P_j - P_i)^{1+\epsilon}}{\partial P_k}$$
$$\frac{\partial (P_j - P_i)^{1+\epsilon}}{\partial P_k} = \begin{cases} (1+\epsilon)(P_j - P_i)^{\epsilon}, & k=j\\ -(1+\epsilon)(P_j - P_i)^{\epsilon}, & k=i\\ 0, & \text{otherwise.} \end{cases}$$

The formulation for  $y_i = -1, y_j \neq -1$  is analogous. Finally, for  $j \neq k$ 

$$\frac{\partial T(P_j)}{\partial P_k} = \begin{cases} -\frac{\delta(1-\delta)(1-\frac{\beta(P_k-\mu_P)(N-1)}{N\sigma_P})}{N(\mu_P-\beta\sigma_P-P_j+(\delta(1-\delta))^{1/\delta})^{\delta}}, & P_j < m \\ \frac{1}{N} \left(\frac{\beta(P_k-\mu_P)(N-1)}{N\sigma_P}-1\right), & m \le P_j \le M \\ \frac{\delta(1-\delta)(1-\frac{\beta(P_k-\mu_P)(N-1)}{N\sigma_P})}{N(P_j-\mu_P-\beta\sigma_P+(\delta(1-\delta))^{1/\delta})^{\delta}} \\ +\frac{2\beta(P_k-\mu_P)(N-1)}{N^2\sigma_P}, & P_j > M, \end{cases}$$

and if j = k,

$$\frac{\partial T(P_j)}{\partial P_j} = \begin{cases} -\frac{\delta(1-\delta)(1-N-\frac{\beta(P_j-\mu_P)(N-1)}{N\sigma_P})}{N(\mu_P-\beta\sigma_P-P_j+(\delta(1-\delta))^{1/\delta})^{\delta}}, & P_j < m\\ 1-\frac{1}{N}\left(\frac{\beta(P_j-\mu_P)(N-1)}{N\sigma_P}-1\right), & M \le P_j \le M\\ \frac{\delta(1-\delta)(N-1-\frac{\beta(P_j-\mu_P)(N-1)}{N\sigma_P})}{\frac{N(P_j-\mu_P-\beta\sigma_P+\delta(1-\delta))^{\delta}}{P_j}} \\ +\frac{2\beta(P_j-\mu_P)(N-n_j)}{N^2\sigma_P}, & P_j > M. \end{cases}$$

We can thus evaluate the derivative of  $\lambda_2(\boldsymbol{\theta})$  with respect to  $\boldsymbol{\theta}$  provided it is simple. We use the non-smooth optimisation method described in [6] to find locally optimal solutions, which alternates between a naive application of gradient descent, in which the simplicity of  $\lambda_2(\boldsymbol{\theta})$  is assumed to hold everywhere, and a descent step based on the directional derivative of  $\lambda_2(\boldsymbol{\theta})$  when it is not simple. We found that the directional step was not required in any of our experiments, and so omit its formulation. Interested readers are referred to the paper [6].

While it is perhaps counterintuitive to increase the similarity between data known to belong to different classes, as in (10), this formulation ensures that projections which do not admit a separation of the known classes are penalised, while those which do admit such a separation allow for partitions which do not include any of the penalised similarities in the ratio cut computation. Figure 1 illustrates this fact. The following lemma shows that projections admitting a separation of the classes have a lower spectral connectivity than those which do not, for small values of the scaling parameter  $\sigma$ .

Lemma 1: Let k be non-increasing, Lipschitz and satisfy  $k(x) \in o(x^{-(1+\epsilon)})$  as  $x \to \infty$ . Let  $\boldsymbol{\theta}_1$  be such that  $\min\{P(\theta_1)_i | y_i = +1\} > \max\{P(\theta_1)_j | y_j = -1\}.$  Then  $\exists \sigma' > 0$  s.t. for any  $0 < \sigma < \sigma'$  and  $\theta_2$  s.t.  $\min\{P(\theta_2)_i | y_i = 0\}$  $+1\} \leq \max\{P(\boldsymbol{\theta}_2)_j | y_j = -1\}$  we have  $\lambda_2(\boldsymbol{\theta}_1) < \lambda_2(\boldsymbol{\theta}_2)$ .

We see then that the formulation given in (10) does indeed induce a penalty, and the penalty forces the optimal solution to admit a separation of the classes if  $\sigma$  is small enough, assuming that the classes can be separated. We can extend the above result to show that the optimal projection converges to the vector normal to the TSVM solution, as  $\sigma \to 0^+$ . We discuss the result in the context of a constrained solution, i.e., one which induces a balanced partition by intersecting a scaled covariance ellipsoid. The result holds for all values of  $\beta$ , and so setting  $\beta$  arbitrarily large proves the result relative to the original TSVM problem.

Lemma 2: Suppose  $\exists v \in \mathbb{R}^d, b \in \mathbb{R}$  s.t.  $y_i(v \cdot x_i - b) >$ 0 for all  $i \in \{1, \dots, l\}$ . Let  $k : \mathbb{R}^+ \to \mathbb{R}^+$  satisfy the following:

- 1) k is non-increasing
- 2) k is Lipschitz
- 3)  $\lim_{x\to\infty} k(x+\epsilon)/k(x) = 0$  for all  $\epsilon > 0$ 1 4)  $k(x) \in o(x^{-(1+\epsilon)})$  as  $x \to \infty$ .

For  $\sigma, \delta > 0$  define  $\boldsymbol{\theta}_{\sigma, \delta} = \operatorname{argmin}_{\boldsymbol{\theta} \in \Theta} \lambda_2(\boldsymbol{\theta}, \sigma, \delta)$ , where  $\lambda_2(\pmb{\theta},\sigma,\delta)$  is the same as  $\lambda_2(\pmb{\theta})$  from before but with an explicit dependence on  $\sigma$  and  $\delta$ . Let  $(v^*, b^*)$  define the largest
Fig. 1. Two projections, one admitting a separation of the classes (Left) and the other not (Right).



+'s and -'s indicate labelled data, while unlabelled data are indicated by o's. The horizontal arrow represents the projection direction. Vertical arrows indicate the maximum projected datum from class -1, say  $p^-$ , and the minimum projected datum from class +,  $p^+$ . The red and blue lines indicate the penalties induced by these two projections. Each unlabelled datum is connected to either  $p^-$  or  $p^+$  by the maximum of these two lines. In the right panel, the minimum of the maximum of these two lines is indicated by the horizontal dashed line, say with value  $\alpha > 0$ . The points  $p^-$  and  $p^+$  are also connected by at least this value, and are connected to their respective classes with similarity H. There is therefore a chain connecting all data with minimum similarity min $\{\alpha, H\}$ . In the left panel, partitioning the data above/below the vertical dashed line leads to a ratio cut with no penalties included.

margin hyperplane which correctly classifies all labelled data and satisfies  $\bar{m} < b^* < \bar{M}$ , where  $\bar{m}$  lies halfway between  $\mu_{v^*\cdot X} - \beta \sigma_{v^*\cdot X}$  and the smallest element of  $P(\theta^*)$  above  $\mu_{v^*\cdot X} - \beta \sigma_{v^*\cdot X}$  and similarly  $\bar{M}$  lies halfway between  $\mu_{v^*\cdot X} + \beta \sigma_{v^*\cdot X}$  and the largest element of  $P(\theta^*)$  below  $\mu_{v^*\cdot X} + \beta \sigma_{v^*\cdot X}$ . Then,

$$\lim_{\sigma,\delta\to 0^+} v(\boldsymbol{\theta}_{\sigma,\delta}) = v^*$$

The distinction between  $\bar{m}$  and  $\mu_{v^{\star}\cdot X} - \beta \sigma_{v^{\star}\cdot X}$ , and analogously for  $\bar{M}$ , is not of much practical concern, but is important for proving the associated theory. The reader may refer [6] for additional discussion.

#### **IV. EXPERIMENTS**

In this section we evaluate the performance of the proposed method, which we will refer to as Semi-Supvervised Spectral Connectivity Projection Pursuit (S<sup>3</sup>CP<sup>2</sup>), on classification data sets taken from the UCI Machine Learning Repository (UCIMLR).<sup>1</sup> and benchmark data sets for semisupervised classification.<sup>2</sup> For each UCIMLR data set we generated 30 sets of labels, 10 for each of the sizes 2%, 10% and 25% the total number of data. The label sets were generated uniformly at random, however sets which did not contain at least one label from each class were rejected and replaced with another. For the data sets taken from [1] we used the same 24 label sets, 12 for each of 10 and 100 labels. The UCIMLR data sets considered were: Mammographic Mass (Mam.): Distinguish benign from malignant masses found during mammography screening. Voters (Vote.): Determine political party affiliation from votes made by US congress members. Breast Cancer (Canc.): Distinguish benign from malignant tumour masses, given physical characteristics. Ionosphere (Iono.): Distinguish radio signals which

show evidence of structure in the ionosphere from those which don't. Parkinsons (Park.): A range of biomedical voice measurements used to determine the presence of Parkinson's disease. The data sets taken from [1] were: g241c: A mixture of 2 Gaussian distributions for which the cluster assumption holds. g241d: A 4 component Gaussian mixture in which cluster structure is misleading for class membership. Digit1: Artificially generated images of the digit "1" varied by translation, ratation, line thickness and length, obscured to increase difficulty. Class boundary at the 0 rotation angle. The cluster assumption holds. USPS: U.S. Postal Service handwritten digits. Digits "2" and "5" form class +1 with the remaining digits forming class -1. The images are obscured using the same transformation as for Digit1. BCI: A braincomputer interface experiment in which a subject imagined movements with their right, class +1, and left hand, class -1. The features are EEG readings taken during the experiment.

We compare performance with a standard SVM<sup>3</sup> trained only on the labelled data, and Semi-Supervised SVM  $(S^{3}VM)^{4}$ . We use the linear kernel for SVM and  $S^{3}VM$ since this provides the most meaningful comparison with our proposed method. Non-linear separators are possible for our method by an explicit embedding of the data within the kernel space, as in [3]. For SVM we use the default parameters given in the package. For S<sup>3</sup>VM we initialised using SVM and built classifiers for each of 5 values of  $C^*$ , which determines the penalty for the unlabelled data violating the large margin, in each experiment using the UCIMLR data sets. We then report the highest average performance of those 5 for each case. For the data sets arising from [1] we use the results presented there. For  $S^3CP^2$  we initialise the projection pursuit using the SVM solution. For the UCIMLR data sets we set  $\beta = 1.5$ ,  $\sigma = 0.3 n n_{0.99} / \sqrt{d}$  where  $n n_{0.99}$  is the 99th centile of the nearest neighbour distances in the data and dis the dimension, and H = u/l, the ratio of the number of

<sup>&</sup>lt;sup>1</sup>M. Lichman. UCI Machine Learning Respository http://archive.ics.uci.edu/ml. Irvine, CA. University of California, School of Information and Computer Science. 2013

<sup>&</sup>lt;sup>2</sup>A selection of data sets used in [1] http://olivier.chapelle.cc/ssl-book/benchmarks.html.

<sup>&</sup>lt;sup>3</sup>We use the R package e1071, which implements the libSVM library <sup>4</sup>We use the SVM-light implementation of T. Joachims available at http://svmlight.joachims.org/

unlabelled and labelled data. The data sets taken from [1] are more challenging, and we selected  $\sigma$  and  $\beta$  from the sets  $\{0.1nn_{0.99}/\sqrt{d}, 0.5nn_{0.99}/\sqrt{d}\}$  and  $\{0.5, 1.5\}$  respectively using cross validation.

Tables I and II report the average classification accuracy over the different label sets. The highest average performance is highlighted in bold.

TABLE I UCIMLR CLASSIFICATION DATA SETS. AVERAGE ACCURACY (%) OVER 10 SPLITS.

	Mam.	Vote.	Canc.	Iono.	Park.			
	2	2% Labelled Examples						
$S^3 CP^2$	79.62	84.53	96.18	66.44	74.71			
SVM	80.58	84.48	91.59	70.82	74.14			
$S^3VM$	79.91	86.90	95.84	72.83	65.50			
	10% Labelled Examples							
$S^3 CP^2$	81.64	90.74	96.04	85.71	79.94			
SVM	80.77	89.00	95.26	80.29	80.23			
$S^3VM$	81.61	89.74	96.66	85.62	71.09			
	25	5% Lat	elled E	Example	es			
$S^3 CP^2$	82.54	90.34	96.49	87.18	80.68			
SVM	82.22	89.33	96.26	84.98	82.60			
S <sup>3</sup> VM	83.04	89.20	96.45	86.92	74.93			

TABLE II SSC Benchmark Data Sets. Average Accuracy (%) over 12 Splits.

	g241c	g241d	Digit1	USPS	BCI			
		10 Labelled Examples						
$S^3 CP^2$	82.02	50.62	89.51	76.17	50.90			
SVM	55.28	56.16	72.90	79.12	52.61			
S <sup>3</sup> VM	79.05	53.65	79.41	69.34	49.96			
	100 Labelled Examples							
$S^3 CP^2$	86.15	72.93	92.77	86.62	70.86			
SVM	74.67	71.98	90.10	86.89	71.56			
S <sup>3</sup> VM	81.82	76.24	81.95	78.88	57.33			

The proposed method achieves the highest average performance in roughly half of the cases considered, and is competitive with the highest performing method in almost all. Importantly the method achieved higher performance than linear  $S^3VM$  in the majority of applications.

Parameter tuning is a difficult task in semi-supervised classification [1], and though our method contains numerous parameters, the majority do not play a significant role in its performance. The parameter  $\sigma$ , and to a lesser extent  $\beta$ , plays the most crucial role in the performance of the method, and determining an appropriate value is necessary for the successful application of the method. We used a simple reference rule which has worked well on many of the examples considered, but believe considerable improvements can be made if a more principled tuning method is employed.

#### V. CONCLUSIONS

We propose a new method for semi-supervised classification, which is based on learning the optimal univariate subspace to perform a binary partition of the projected data set using semi-supervised spectral clustering. The labels of the training data are incorporated into the model in such a way that the globally optimal solution must admit a separation of the classes within the training data, if such a solution exists, and for all scaling parameters close to zero. We also show that asymptotically this optimal solution converges to the subspace normal to the optimal TSVM hyperplane, as the scaling parameter is reduced to zero, thereby providing a theoretical connection between our proposed method and popular semi-supervised classification methodology. Experimental results indicate the proposed method is competitive with state-of-the-art TSVM implementation in terms of classification accuracy.

### **APPENDIX: PROOFS**

The following lemma is useful for proving Lemmas 1 and 2.

Lemma 3: Let k be non-increasing and Lipschitz with constant K, and let k(0) = 1. For  $\theta \in \Theta$  let

$$\Delta_{\boldsymbol{\theta}} = [\bar{m}_{\boldsymbol{\theta}}, \bar{M}_{\boldsymbol{\theta}}] \cap [\max\{P(\boldsymbol{\theta})_i | y_i = -1\}, \min\{P(\boldsymbol{\theta})_j | y_j = +1\}],$$

where

$$\bar{m}_{\boldsymbol{\theta}} = \frac{\mu_{P(\boldsymbol{\theta})} - \beta \sigma_{P(\boldsymbol{\theta})} + \min\{P(\boldsymbol{\theta}) \cap [\mu_{P(\boldsymbol{\theta})} - \beta \sigma_{P(\boldsymbol{\theta})}, \infty)\}}{2}$$
$$\bar{M}_{\boldsymbol{\theta}} = \frac{\mu_{P(\boldsymbol{\theta})} + \beta \sigma_{P(\boldsymbol{\theta})} + \max\{P(\boldsymbol{\theta}) \cap (-\infty, \mu_{P(\boldsymbol{\theta})} + \beta \sigma_{P(\boldsymbol{\theta})})]\}}{2}$$

Let  $\sigma' = \frac{K}{(1+\epsilon)^{1/1+\epsilon}}$ . If  $\Delta_{\theta} \neq \emptyset$  then set  $G_{\theta} = \max_{b \in \Delta_{\theta}} \min_{i \in \{1, \dots, N\}} |P(\theta)_i - b|$ . Then for  $\sigma \leq \sigma'$  we have,

$$\lambda_2(\boldsymbol{\theta}) \ge \min\left\{\frac{1}{9|X|^3}k\left(\frac{2G_{\boldsymbol{\theta}} + \delta D}{\sigma}\right), \frac{1}{9|X|}\left(\frac{\sigma}{K}\right)^{1+\epsilon}\right\},$$

where  $D = \max{\{\text{Diam}(X), \text{Diam}(X)^{1-\delta}\}}$ . If  $\Delta_{\theta} = \emptyset$  then we simply have

$$\lambda_2(\boldsymbol{\theta}) \ge \frac{1}{9|X|} \left(\frac{\sigma}{K}\right)^{1+\epsilon}$$

for all  $\sigma \leq \sigma'$ .

Proof: Since k has Lipschitz constant K we have  $k(x/\sigma) + x^{1+\epsilon} \ge (k(0) - \frac{K}{\sigma}x)^+ + x^{1+\epsilon} = (1 - \frac{K}{\sigma}x)^+ + x^{1+\epsilon}.$ Since  $\sigma \le \sigma' = \frac{K}{(1+\epsilon)^{1/1+\epsilon}}$  we can show that  $(1 - \frac{K}{\sigma}x)^+ + x^{1+\epsilon} \ge (\frac{\sigma}{K})^{1+\epsilon}$  for all  $x \ge 0.$ 

First consider the case  $\Delta_{\boldsymbol{\theta}} = \emptyset$ . This implies that for each j, either  $P(\boldsymbol{\theta})_j \leq \max\{P(\boldsymbol{\theta})_i | y_i = -1\}$  or  $P(\boldsymbol{\theta})_j \geq \min\{P(\boldsymbol{\theta})_i | y_i = +1\}$ . Let I be the index corresponding to  $\max\{P(\boldsymbol{\theta})_i | y_i = +1\}$ . Therefore, for each j s.t.  $x_j \in \mathbf{X}^U$  either  $A_{jI} \geq k(|P(\boldsymbol{\theta})_j - P(\boldsymbol{\theta})_I)|/\sigma) + |P(\boldsymbol{\theta})_j - P(\boldsymbol{\theta})_I|^{1+\epsilon} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}$  or  $A_{jJ} \geq k(|P(\boldsymbol{\theta})_j - P(\boldsymbol{\theta})_J|/\sigma) + |P(\boldsymbol{\theta})_j - P(\boldsymbol{\theta})_J|^{1+\epsilon} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}$ . In addition we have  $A_{IJ} \geq k(|P(\boldsymbol{\theta})_I - P(\boldsymbol{\theta})_J|/\sigma) + |P(\boldsymbol{\theta})_J - P(\boldsymbol{\theta})_J|/\sigma) + |P(\boldsymbol{\theta})_I - P(\boldsymbol{\theta})_J|/\sigma) + |P(\boldsymbol{\theta})_I - P(\boldsymbol{\theta})_J|^{1+\epsilon} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}$ . and for each j s.t.  $y_j = +1$  and i s.t.  $y_i = -1$  we similarly have  $A_{jI}, A_{iJ} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}$ . Now, let u be the second eigenvector of  $L(\theta)$ , then  $||u|| = 1, u \perp 1$  and so  $\exists i, j$  s.t.  $u_i - u_j \geq \frac{1}{\sqrt{|X|}}$ . If  $|u_I - u_J| \leq \frac{1}{3\sqrt{|X|}}$  then either  $|u_i - u_I| \geq \frac{1}{3\sqrt{|X|}}$  and  $|u_i - u_J| \geq \frac{1}{3\sqrt{|X|}}$  or  $|u_j - u_I| \geq \frac{1}{3\sqrt{|X|}}$  and  $|u_j - u_J| \geq \frac{1}{3\sqrt{|X|}}$ . Then,  $\lambda_2(L(\theta)) = u \cdot L(\theta)u = \frac{1}{2}\sum_{k,l}A_{kl}(u_k - u_l)^2 \geq A_{iI}(u_i - u_I)^2 + A_{iJ}(u_i - u_J)^2 + A_{jI}(u_j - u_I)^2 + A_{jJ}(u_j - u_J)^2 \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon} / 9|X|$  in all possible cases, since  $\left(\frac{\sigma}{K}\right)^{1+\epsilon} \leq 1 \leq H$ . On the other hand we have  $|u_I - u_J| \geq \frac{1}{3\sqrt{|X|}}$  and so

$$\begin{split} \lambda_2(L(\pmb{\theta})) &\geq A_{IJ}(u_I - u_J)^2 \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}/9|X| \text{ as required.}\\ \text{Now consider the case } \Delta_{\pmb{\theta}} \neq \emptyset. \text{ Define } u, I, J \text{ as above.}\\ \text{If } \max_{i,j:P(\pmb{\theta})_i,P(\pmb{\theta})_j\in[P(\pmb{\theta})_I,P(\pmb{\theta})_J]}(u_i - u_j) &\leq \frac{1}{3\sqrt{|X|}}, \text{ then}\\ \text{since } \exists i, j \text{ with } u_i - u_j \geq \frac{1}{\sqrt{|X|}} \text{ we must have either } P(\pmb{\theta})_i \notin [P(\pmb{\theta})_I, P(\pmb{\theta})_J] \text{ and } |u_i - u_I| \geq \frac{1}{3\sqrt{|X|}} \text{ and } |u_i - u_J| \geq \frac{1}{3\sqrt{|X|}} \text{ or } P(\pmb{\theta})_j \notin [P(\pmb{\theta})_I, P(\pmb{\theta})_J] \text{ and } |u_j - u_I| \geq \frac{1}{3\sqrt{|X|}} \text{ and } |u_j - u_J| \geq \frac{1}{3\sqrt{|X|}} \text{ and } |u_j - u_J| \geq \frac{1}{3\sqrt{|X|}}. \text{ Suppose w/o loss of generality that } P(\pmb{\theta})_i \text{ satisfies these three conditions. If } x_i \in \mathbf{X}^U \text{ then since } P(\pmb{\theta})_i \notin [P(\pmb{\theta})_I, P(\pmb{\theta})_J] \text{ we have either } A_{iI} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon} \text{ or } A_{iJ} \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon} \text{ and the result follows similar to above. If } x_i \in \mathbf{X}^L \text{ then either } A_{iI} = 1 \text{ or } A_{iJ} = 1 \geq \left(\frac{\sigma}{K}\right)^{1+\epsilon}, \text{ and again the result follows as above. If instead we have } \max_{i,j:P(\pmb{\theta})_i,P(\pmb{\theta})_j\in[P(\pmb{\theta})_I,P(\pmb{\theta})_J]}(u_i - u_j) > \frac{1}{3\sqrt{|X|}}, \text{ then the result follows analogously to [6, Lemma 2], with the addition of the factor <math>\frac{1}{3}$$
 on the distance between elements of u in consideration.

Proof of Lemma 1 Proof: Let  $G_{\theta_1} > 0$  be defined as in Lemma 3, and let b be the corresponding point at which the distance is maximised. Let L and R be the number of projected data lying to the left and right of b respectively. Then, since spectral clustering solves the relaxation of the ratio cut, we have

$$\lambda_{2}(\boldsymbol{\theta}_{1}) \leq \frac{1}{|X|} \min_{\substack{C \subset X \\ x_{j} \notin C}} \sum_{\substack{i,j:x_{i} \in C \\ x_{j} \notin C}} A(\boldsymbol{\theta}_{1})_{ij} \left(\frac{1}{|C|} + \frac{1}{|X \setminus C|}\right)$$
$$\leq \frac{1}{|X|} \sum_{\substack{i,j:P(\boldsymbol{\theta}_{1})_{i} < b \\ P(\boldsymbol{\theta}_{1})_{j} > b}} k(d(P(\boldsymbol{\theta}_{1})_{i}, P(\boldsymbol{\theta}_{1})_{j})) \left(\frac{1}{L} + \frac{1}{R}\right)$$
$$\leq k(2G\boldsymbol{\theta}_{i}, /\sigma).$$

For any  $\theta_2$  s.t.  $\Delta_{\theta_2}$  defined as in Lemma 3 is empty, we have

$$\frac{\lambda_2(\boldsymbol{\theta}_1)}{\lambda_2(\boldsymbol{\theta}_2)} \le 9|X|K^{1+\epsilon}k(2G_{\boldsymbol{\theta}_1}/\sigma)\sigma^{-(1+\epsilon)}$$

This right hand side is independent of  $\theta_2$ , and converges to 0 as  $\sigma \to 0^+$  since  $k(x) \in o(x^{-(1+\epsilon)})$  as  $x \to \infty$ . Therefore the result holds.

Proof of Lemma 2 Proof: By Lemma 1 we know that  $\exists \sigma' > 0$  s.t.  $0 < \sigma < \sigma' \Rightarrow \Delta_{\boldsymbol{\theta}_{\sigma,\delta}} \neq \emptyset$ , where  $\Delta_{\boldsymbol{\theta}}$  is as in Lemma 3. Take  $\gamma > 0$ . It has been shown [9] that  $\exists m_{\gamma} > 0$  s.t. for  $w \in \mathbb{R}^d$  and  $c \in \mathbb{R} ||(w,c)/||w|| - (v^*, b^*)|| > \gamma \Rightarrow \max(w/||w||, c/||w||) < \max(v^*, b^*) - m_{\gamma}$ . Let  $\boldsymbol{\theta}^*$  be

such that  $v(\theta^*) = v^*$  so that  $\operatorname{margin}(v^*, b^*) = G_{\theta^*}$ , where  $G_{\theta}$  is as in Lemma 3. As in the proof of Lemma 1 we have

$$\lambda_2(\boldsymbol{\theta}^\star) \le k(2G_{\boldsymbol{\theta}^\star}/\sigma)$$

In addition, for small enough  $\sigma>0$  we have

$$\frac{1}{9|X|^3}k\left(\frac{2G+\delta D}{\sigma}\right) < \frac{1}{9|X|}\left(\frac{\sigma}{K}\right)^{1+\epsilon}$$

holding uniformly in  $\delta > 0$  for any G > 0, where K, D are as in Lemma 3. Therefore, for small  $\sigma, \delta$  we have

$$\lambda_2(\boldsymbol{\theta}_{\sigma,\delta}) \geq \frac{1}{9|X|^3} k\left(\frac{2G_{\boldsymbol{\theta}_{\sigma,\delta}} + \delta D}{\sigma}\right)$$

and hence

=

$$\frac{1}{9|X|^3}k\left(\frac{2G_{\pmb{\theta}_{\sigma,\delta}}+\delta D}{\sigma}\right) \leq k\left(\frac{2G_{\pmb{\theta}^\star}}{\sigma}\right)$$

since  $\lambda_2(\boldsymbol{\theta}_{\sigma,\delta}) \leq \lambda_2(\boldsymbol{\theta}^{\star})$ . Now, take  $\delta'$  s.t.  $\delta'D < \frac{m_{\gamma}}{2}$ . Since  $\lim_{x\to\infty} k(x+\epsilon)/k(x) = 0$  for all  $\epsilon > 0$ ,  $\exists \sigma' > 0$  s.t.  $9|X|^3K(2G_{\boldsymbol{\theta}^{\star}}/\sigma) < k((2G_{\boldsymbol{\theta}^{\star}} - \frac{m_{\gamma}}{2})/\sigma)$  for all  $\sigma < \sigma'$ . For  $\sigma < \sigma', \delta < \delta'$  we have,

$$\begin{split} k\left(\frac{2G_{\boldsymbol{\theta}_{\sigma,\delta}} + \frac{m_{\gamma}}{2}}{\sigma}\right) &\leq k\left(\frac{2G_{\boldsymbol{\theta}_{\sigma,\delta}} + \delta D}{\sigma}\right) \\ &\leq 9|X|^{3}k\left(\frac{2G_{\boldsymbol{\theta}^{\star}}}{\sigma}\right) \\ &\leq k\left(\frac{2G_{\boldsymbol{\theta}^{\star}} - \frac{m_{\gamma}}{2}}{\sigma}\right) \\ &\Rightarrow 2G_{\boldsymbol{\theta}_{\sigma,\delta}} + \frac{m_{\gamma}}{2} &\geq 2G_{\boldsymbol{\theta}^{\star}} - \frac{m_{\gamma}}{2} \\ &\Rightarrow \max_{b \in \Delta \boldsymbol{\theta}_{\sigma,\delta}} \max(v(\boldsymbol{\theta}_{\sigma,\delta}), b) &\geq \max(v^{\star}, b^{\star}) - \frac{m_{\gamma}}{2} \\ &\Rightarrow \|(v(\boldsymbol{\theta}_{\sigma,\delta}), b) - (v^{\star}, b^{\star})\| &\leq \gamma \\ &\Rightarrow \|v(\boldsymbol{\theta}_{\sigma,\delta}) - v^{\star}\| &\leq \gamma \end{split}$$

Since  $\gamma > 0$  was arbitrary and the above holds for all  $\sigma < \sigma', \delta < \delta'$  we must have  $\lim_{\sigma, \delta \to 0^+} v(\theta_{\sigma, \delta}) = v^*$  as required.

#### REFERENCES

- O. Chapelle, B. Schölkopf, A. Zien (eds.) Semi-Supervised Learning MIT Press, Cambridge, 2006.
- [2] V. Vapnik and A. Sterin. On structural risk minimization or overall risk in a problem of pattern recognition. Automation and Remote Control, 10(3):14951503, 1977.
- [3] O. Chapelle, M. Chi, and A. Zien, A. A continuation method for semisupervised SVMs. Proc ICML, 2006.
- [4] T. Joachims. Transductive inference for text classification using support vector machines. Proc. ICML, 1999.
- [5] U. Von Luxburg. A Tutorial on Spectral Clustering. Statistics and Computing, 17:395-416, 2007.
- [6] D. Hofmeyr, N. Pavlidis and I. Eckley. Minimum Spectral Connectivity Projection Pursuit for Unsupervised Classification. arXiv preprint, arXiv:1509.01546, 2015.
- [7] D. Wagner and F. Wagner. Between Min Cut and Graph Bisection. Springer, 1993.
- [8] S. Guattery and G. Miller. On the Quality of Spectral Separators. SIAM Journal of Matrix Analysis and Applications, 19(3):701-719, 1998.
- [9] N. Pavlidis, D. Hofmeyr and S. Tasoulis. Minimum Density Hyperplane: An Unsupervised and Semi-supervised Classifier. arXiv preprint, arXiv:1507.04201, 2015.
- [10] O. Chapelle and A. Zien. Semi-supervised classification by low density separation. Proc. AISTATS, 2005.

# Face and Iris biometrics person identification using hybrid fusion at feature and score-level

Valentine Azom

School of Mathematics, Statistics and Computer Science University of KwaZulu-Natal Email: 215041656@stu.ukzn.ac.za

Aderemi Adewumi

School of Mathematics, Statistics and Computer Science University of KwaZulu-Natal Email: adewumia@ukzn.ac.za

Jules-Raymond Tapamo Center for Applied Artificial Intelligence Center for Applied Artificial Intelligence Center for Applied Artificial Intelligence School of Engineering University of KwaZulu-Natal Email: tapamoj@ukzn.ac.za

Abstract—Face and Iris biometrics are amongst the studied unimodal systems by researchers over the past years due to their ease of acquisition and accuracy, respectively, during the recognition process. However unimodal systems are not perfect when deployed to real-world applications due to non-ideal conditions such as off-angle gaze, illumination, occlusion and variation in posing. These limitations have led to an increase of research in multi-biometrics. In recent times, researchers have worked on combining unimodal templates with different methods with each having to compensate for the shortcomings of the unimodal systems. In this work we present a hybridized fusion strategy that combines, three classifiers based on feature and score level fusion using a decision level fusion rule. We compare the recognition rate of the proposed method with other fusion methods in literature. We have obtained a recognition accuracy of 98.75%. The proposed method was validated using the ORL face and CASIA iris datasets.

Keywords—Multimodal biometrics; Face recognition; Iris recognition; Feature extraction; Score fusion, feature fusion and decision fusion

# I. INTRODUCTION

Biometrics is the process of using human traits such as the iris, face, hand geometry, vein and fingerprint for recognition purposes and this has attracted interest in security technology. Biometric applications have been deployed in various sectors of the economy especially in high risk security areas, with emphasis placed on the accuracy of the system. Face and iris biometrics have become popular over the years but a couple of setbacks have made these unimodal systems less effective for real-world applications, Common problems encountered include: varying illumination conditions, off-angle gaze, occlusion, noise and spoof attacks. Multi-modal biometrics has been identified as one way to overcome most of the shortcomings faced by unimodal systems.

multi-biometric systems can be classified into four categories namely: 1) multi-sample that collate different samples of the same biometric trait; 2) multi-instance that collate different instances of traits e.g. using right and left iris images for recognition; 3) multi-algorithm that use two or more feature extraction methods for recognition; 4) multi-modal that combine biometric traits from two or more modalities for recognition [1].

Moreover, the fusion process for multi-biometric systems can be categorized as taking place in two levels namely: 1) before matching 2) after matching, as seen in figure 1. Fusion methods that take place prior to matching are sensor and feature based. In sensor based fusion, images are captured with one or more sensors and are later stitched together before matching. In feature level of fusion, features are extracted from each modality and are concatenated either in series or parallel to form a single feature vector, leading to a feature vector with high dimensionality. In this case feature selection methods are applied to reduce the dimensionality of the feature vector. Post matching fusion methods include the score, rank and decision level fusion [4].

In score level fusion, matching scores are obtained from different modalities and can be combined based on transformation, classification and density methods. In transformation methods, matching scores obtained from different modalities are changed into a single domain using normalization methods like the Z-score, Min-Max and Tanh methods. For classifier based methods matching scores are passed as features to be trained by a classifier that detects either user is a genuine user or an imposter. Density based methods use statistical distribution estimation methods for combining scores to then make the final decision. Both rank and decision level fusion are under studied in literature due the loss of information during the recognition process. However decision level of fusion involves combining output labels from different modalities to make a final decision either using Boolean operations or voting rules [4].

In this work we propose a hybrid fusion strategy both at feature and score level. For each of the face and iris modalities, five feature extraction algorithms has been applied as local and global feature extractors, namely Local Binary Pattern Histogram (LBPH), modular Principal Component Analysis (mPCA) and sub-pattern Principal Component Analysis (sp-PCA) as local methods and Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA) as global methods [2]. Feature level fusion is then performed for each modality on the feature vectors obtained from the five methods to form the first two classifiers. Next we perform weighted score-level fusion using matching scores obtained from LDA



Fig. 1: Levels of fusion in multi-modal systems, FE= feature extraction, FU =fusion unit, MM=Matching module, DM =Decision module, A/R= Accept/Reject [4]



Fig. 2: Proposed hybrid fusion scheme, FU= Feature-level fusion, SU=Score-level fusion DU=Decision-level fusion.

for face and LBPH for Iris to form the third classifier. Outputs obtained from each of the classifiers is then fused at decision level using the majority voting rule to obtain the final decision (genuine or imposter user). A diagrammatic representation is shown in figure 2.

The rest of the paper is organized as follows: Section II discusses related work, the local and global feature extraction methods described in Section III, Section IV discusses face recognition, Section V we discuss the iris recognition process, Section VI discusses normalization methods, Section VII feature fusion for face and Iris vectors, Section VIII

Score level fusion of global and local features, Section IX proposed system, Section X we discuss about the experiments and results obtained and in Section XI we conclude the paper.

# II. RELATED WORK

In this section we discuss the work done by researchers mainly on score-level fusion methods, later delving into other fusion methods. A multimodal biometric study on Facial and Iris capture is mentioned in [1]. A multi-sensor approach was used for both face and iris image acquisition. The videos obtained from the three different Iris on the Move (IOM) sensors were stitched together to obtain a single unit. Both Viola Jones [3] and modified Daugman's algorithm [2] were used to extract features from the respective biometric templates. Matching scores obtained from the individual modalities were normalized and a weighted sum rule were used as the fusion technique.

Maryam et al [4], proposed a new approach for score-level fusion based on face and iris matching scores. They used both local and global features to obtain matching scores from each modality. Then they fused the matching scores obtained from face and iris templates separately and concatenated them into a single score vector. The Nearest neighbor classifier was used for identification with recognition rate of 98.25%.

A multimodal biometric system was proposed in [5]. This was based on combining unimodal templates from both voice and fingerprint. A Support Vector Machine (SVM) was used as a classifier rule for binary verification decision making, combining test scores from the individual modalities. Comparison experiments were validated on FingerCell fingerprint database and ELSDSR voice database. Using rank identification, it was shown that score fusion technique gave 70% and 100% recognition accuracy at the First and Second rank. While, unimodal systems gave 60% (fingerprint), 50% (voice) recognition accuracy at First rank and 100% accuracy at the fourth and eight rank respectively.

Wang et al [6] presented a new score fusion methodology using support vector machines. Face and Iris features were extracted using Laplacian and phase methods. The matching scores obtained from each modality was concatenated into a feature vector. The feature vector was then passed into an SVM classifier to output a fused score which depending on the specified system threshold, would classify the input as an imposter or genuine user. Results obtained from this method was compared with known score fusion methods and it provided better accuracy with Equal Error Rate (EER) of 0.35%.

Yunhong et al [7] combined face and iris biometric traits using three classifier and transformation based fusion rules. They extracted features from the unimodal systems using PCA and 2D-signals and compared the total error rate between the weight sum, fisher discriminant and radial based neutral network methods. The Radial Basis Function Neural Network (RBFN) produced the lowest error rate amongst the three methods.

Qian et al [8] proposed a threshold-optimized "AND" and "OR" rule based decision scheme. They showed that optimizing the threshold values of the classifiers of individual modalities provided substantial improvements to the performance of the fusion system by balancing matching scores from individual classifiers. The benefit of this is that the matching score normalization process performed in other fusion schemes won't be required, thus reducing the risk of performance degradation as component classifiers had significantly different performances. Experimental results showed improvements over original classifiers that were fused with results comparable with other conventional fusion schemes. Works in literature have shown that combination of biometric templates at different fusion levels improves the accuracy of the system. However less work has been done to explore the possibility of combining results at different fusion levels in making the final decision.

# III. FEATURE EXTRACTION

The local feature extraction methods used in this work extract features based on a sub-region defined within the image rather than considering the image as a whole. This operation is repeated across the whole image until all the sub-features are extracted. Global feature extraction on the other hand operate on the image as a whole to extract features for recognition. One of the main advantages of using local feature extractors is that features obtained have lower dimensionality and a reduced error rate due to partial illumination that may occur in some parts of the image. Below a description of the two categories of the feature extraction methods.

# A. Global methods

1) Principal component analysis: This method tries to maximize the total scatter of the centered images in the training set. It reduces the dimensionality of the training set to a one dimension d eigenface vector space. All test images are projected to the eigenfaces space to obtain feature vectors with the same dimension. Matching scores are obtained by finding the Euclidean distance between the feature vector of the test image and feature vectors obtained from the training set [4].

2) Linear discriminant analysis: LDA is closely related to PCA as they both try to project data into a vector space. However, LDA explicitly attempts to model the difference between classes in the data. It maximizes the inter-class variance and minimizes the intra-class variance [4]. Matching scores are obtained by finding Euclidean distance between the feature vector of the test image and feature vectors obtained from the training set.

# B. Local Methods

1) Local binary pattern histogram: The basic concept of LBPH is an 8-bit or 16-bit operator used on a specific area of the image  $(3 \times 3 \text{ or } 9 \times 9)$  such that the pixels greater in value than the center pixel is assigned a value of 1, otherwise 0 is assigned [4]. We have used the circular LBPH with extended neighborhood. A circle with specified radius is defined within every block in the image in which the sampling points (neighbors) are located at the edge of the circle. Matching scores are obtained by finding the Chi-square distance between the feature vector of the test image and feature vectors obtained from the training set.

2) Sub-pattern principal component analysis: This is a variant of PCA in the sense that images are first turned into row vectors and divided into k sub-patterns. PCA is then performed on each sub-pattern group of the training set [9]. Matching scores are obtained by computing the average Euclidean distance between the feature vector of a test image and feature vector of images in the training set from each k sub-pattern.



(a) ORL image divided into (b) A circular ring for LBP  $3 \times 3$  block with 8 neighbors and radius of 1

Fig. 3: face block and LBPH ring

3) Modular principal component analysis: This is also a variant of PCA that divides the image into blocks and then performs PCA on the set of divided images [7]. All divided images are converted into row vectors and the number of images in the training set is increased to  $(N \times d)$ . Where N and d are number of blocks and d is the number images in the training set. Matching scores are obtained by finding the Euclidean distance between feature vector of the test image and feature vectors obtained from the training set [9].

#### IV. FACE RECOGNITION

All five feature extraction methods were applied to the ORL face dataset. Each image in the dataset is of size  $112 \times 92$  with 10 images each for 40 subjects with different pose, facial expression and varying illumination conditions [10]. The eigenvectors corresponding to non-zero eigenvalues are extracted for computing feature vectors of the global methods. For the local methods we have used N=9 partitions. Meaning that each image in the training set, was divided into  $3 \times 3$  blocks. In LBPH method, we have employed the (K,R) circular neighborhood operation where K is the number of neighbors and R is the radius of the circular neighborhood as seen in figure 3. In our case K = 8 and R = 1. Finally we use Euclidean distance in calculating matching scores between the test and training images using eq 1.

$$\sum_{j}^{k} \left\| X - Y_{j} \right\|_{2} \tag{1}$$

Where X is the projection of the probe image and  $Y_j$  is the projection of the  $j^t h$  image in the training set for  $j = 1, 2, \dots, k$ . k is number of training samples.

#### V. IRIS RECOGNITON

Pre-processing is needed prior to extracting features from an iris image. In fact, most iris images are corrupted with noise such as the specular light, light produced by illumination during acquisition. Here we have used CASIAv3 (iris-Interval) iris dataset 249 subjects with left and right iris images [11]. The iris pre-processing is divided into two stages described below:



Fig. 4: Segmented image



Fig. 5: Normalized image

#### A. Segementation

Segmentation is the process of locating the pupil and Iris in the raw data provided from the dataset. This can be quite difficult as the eye image could be corrupted with noise such as the specular lights introduced during acquisition and the presence of eyelashes that occlude the iris part of the eye. Viterbi's algorithm [12] is used extracting the iris and pupil region. After pupil and Iris contour estimation we remove noise due to the occlusion of the eye-lashes by using Full Width at Half Maximum (FWHM) of the Gaussian density function, modelling the histogram of iris pixel intensities [12]. Pixel  $I_{i,j}$  is an occlusion if:

$$\mid I_{i,j} - \mu \mid = 2.35\sigma \tag{2}$$

Where  $\mu$  and  $\sigma$  are respectively the mean and variance of the pixel intensity of the iris image.

#### B. Normalization

The normalization method is based on the Daugman's rubber sheet model [2]. It transforms the area of the iris delimited by the iris and pupil contours detected in the segmentation process into a scale-invariant and pupil-dilation-invariant band.

Coarse pupil and iris contours are detected by Viterbi algorithm by selecting the minimum number of noisy points and angles that define a closed contour [12].

Let L and B represent the length and breadth of the rectangular sheet. we compute  $\theta_i \in [0, 2\pi]$  such that :

$$\theta_i = 2\pi \frac{i}{B}, i = 0, 1 \cdots B \tag{3}$$

Let  $(x_p, y_p, \phi_p)$  and  $(x_r, y_r, \phi_r)$  represent coordinate points of the coarse pupil and iris contours respectively where x and y are coordinate point of the radius relative to the center and  $\phi$ is the angle of non-regular sampling. We estimate  $(X_i^r, Y_i^r)$ corresponding to  $\theta_i$  for two nearest points s, s + 1 as in :

$$X_i^r = (1 - \alpha)x_r^s + \alpha x_r^{s+1} \tag{4}$$

$$Y_i^r = (1 - \alpha)y_r^s + \alpha y_r^{s+1} \tag{5}$$

$$\alpha = \frac{\theta_i - \phi_r^s}{\phi_r^{s+1} - \phi_r^s} \tag{6}$$

#### VI. NORMALIZATION METHODS

Matching scores obtained from the face and iris feature extractors are not homogenous hence the need to transfer them to a common domain. Some normalization techniques have been proposed in literature and they include Min-Max, Tanh and Z-score normalization methods [4]. Min-Max normalization is given by:

$$S'_{k} = \frac{S_{k} - Min}{Max - Min}, \quad where \ k = 0, 1 \cdots k$$
(7)

Where  $S'_K$  and  $S_K$  are the normalized score and corresponding matching score. Tanh normalization is given by:

$$S'_{k} = \left\{ \tanh\left(0.001\left(\frac{S_{k} - \mu_{gh}}{\sigma_{gh}}\right)\right) \right\}$$
(8)

Where  $\mu_{gh}$  and  $\sigma_{gh}$  are the mean and standard deviation of the set of matching scores. Z-score normalization is given by:

$$S'_k = \frac{S_k - \mu}{\sigma} \tag{9}$$

Where  $\mu$  and  $\sigma$  are the mean and standard deviation of the set of matching scores.

#### VII. FEATURE FUSION OF FACE AND IRIS VECTORS

After performing the pre-processing and feature extraction stages for both face and Iris biometric templates, we perform feature-level fusion for each of the face and Iris templates separately. Considering that the magnitude and dimensions of features obtained from each feature extraction methods might not be same. We have employed Z-score normalization technique discussed in the previous section to bring all feature values to be into the same domain. For spPCA and mPCA methods we have divided the each image in the training data into 9 sub-patterns and a  $3 \times 3$  block respectively. Therefore nine feature vectors per image are obtained from this methods, together with single feature vector each from PCA, LDA and LBPH to bring the total number for feature vectors after fusion to twenty one.

Let  $X_i$  and  $Y_i$  represent features extracted from face and Iris respectively; where  $i = 1, 2 \cdots N$  where N is the number of features extracted from each methods. In our case N = 21, as single feature vectors are obtained per image for PCA, LDA and LBPH methods, while nine feature vectors per image are each extracted for spPCA and mPCA methods. Then the fusion of feature vectors  $Z_f$  and  $Z_I$  for face and Iris respectively are expressed as:

 $Z_f = \{X_1, X_2 \cdots X_N\}$  and  $Z_I = \{Y_1, Y_2 \cdots Y_N\}$ 

In order to perform recognition on this method we consider two instances of the fused face vectors  $Z_{1f} = \{X_{11}, X_{12}, \dots, X_{1N}\}$  and  $Z_{2f} = \{X_{21}, X_{22}, \dots, X_{2N}\}$ . We obtain the average Euclidean distance d as:

$$d = \frac{S_1 + S_2 + \dots + S_N}{N}$$
(10)

Where  $S_1 = ||X_{21} - X_{11}||_2$  and  $S_N = ||X_{2N} - X_{1N}||_2$ The process is also performed for the iris template, then nearest neighbour classifier is now used for classification.

# VIII. SCORE LEVEL FUSION OF LOCAL AND GLOBAL METHODS

We have performed score-level of fusion based on subspace LDA for face and LBPH for Iris because of their performance on the individual unimodal systems. We have employed the weighted score fusion for this stage of our work and it is computed as follows [4].

$$WS = W_F \times S_F + W_I \times S_I, where W_F + W_I = 1 \quad (11)$$

Where  $W_F$  and  $W_I$  represents the reliability of each template towards the fusion process. This weights have been determined empirically.

#### IX. PROPOSED METHOD

In our proposed method we have chosen to hybridize feature, score and decision level fusion by combining classifiers based on the prior two fusion strategies. We have employed this process as feature level fusion contains more information about each biometric template and ability of score-level fusion to combine different noise levels from both templates. While decision level fusion has the ability to combine outputs from the two fusion levels [13].

Different decision level techniques have been proposed including maximum, minimum and average voting. In this work we have chosen to use majority voting technique. In majority voting, each classifiers outputs its own class and the class with the highest occurrence amongst all the classifiers is chosen. If there is a tie, then the classifier with the highest matching score is chosen as the identified class as in eq 12. In maximum voting the classifier output with the highest value is chosen as the identified class [13] as in eq 13.

$$H(x) = argmax_{i=1}^{k}(y_i) \tag{12}$$

Where y and x are output and input respectively. In average voting involves finding the mean of the output confidence for each class across all of the ensemble. The class highest mean value is identified as the correct class [13].

$$H(x) = argmax_{i=1}^{k} \left(\frac{1}{k} \sum_{i=1}^{k} y_{ij}(x)\right)$$
(13)

Where y and x are output and input respectively.

### X. EXPERIMENTS AND RESULTS

The performance of the system is tested on a chimeric dataset constructed from ORL Face and CASIA Iris database. The ORL face dataset [8] consist of 10 images each of 40 individuals taken at different face angle positions. We randomly selected seven images and assigned three images for training and four images testing meaning a total of 120 images was used for training and 160 images was used for testing . The same procedure has been applied on the iris images in the CASIA database [9]. We have calculated the recognition accuracy using the formula below:

$$RR(\%) = \frac{No \ of \ correctly \ predicted \ labels}{No \ of \ test \ labels} \times 100 \quad (14)$$

# A. Uni-modal systems

The first set of empirical procedures performed was to obtain the recognition rate for both face and Iris biometrics individually using the local and global feature extractors. In general, results obtained from the experiment show that LDA performed the best for face recognition and LBP performed also best for Iris recognition as seen in table I.

Method	ORL %	CASIA %
PCA	87.50	85.63
LDA	90.25	87.50
LBPH	87.50	92.38
spPCA	88.35	86.38
mPCA)	88.20	91.50

TABLE I: Recognition rate for uni-modal systems

# B. Proposed system

In order to test the effectiveness of the proposed system we will compare the recognition rate obtained with it from the individual classifiers and other fusion methods in literature. Results obtained show that proposed method out performs other multi-biometric classifiers shown in table II.

Method	Performance %
Face fused vector	90.25
iris fused vector	93.25
weighted score-fusion	97.50
Proposed-FU	<b>98.75</b>

TABLE II: Recognition rate for multi-modal systems

The proposed method performes better than many other feature or score-level fusion techniques because it combines information from three different classifiers using the voting method to make the final decision. While other methods make decision based on either feature-level fusion or score-level fusion. The proposed system takes advantage of the rich information inherent in feature fusion and ability to combine modalities based on noise level through weighted score fusion to improve the overall performance of the system. Table III shows a comparison of the proposed method with other methods in literature and it can be seen that the proposed method performs better than other methods.

Method	Performance %
SumBoth [1]	93.2
BordaBoth-Exponential [1]	92.0
Score rule with SVM [4]	98.25
Proposed method	98.75

TABLE III: Comparison of the proposed methods with other fusion methods

# XI. CONCLUSION

In this work we have implemented a hybrid fusion scheme based on feature, score and decision level of fusion by combining three fusion classifiers with a decision rule. First we used both global and local feature extraction methods on both face and Iris biometric template. Then we performed feature-level fusion individually for face and Iris to form the fused face and Iris classifiers. Followed weighted score fusion between LDA method for face and LBPH method for Iris as both produced the highest recognition rate for the individual modalities. We combined three classifiers with a popular decision level fusion rule in order to obtain the final decision. As shown previously, results obtained from the proposed strategy outperform other fusion systems in literature.

Further work will concentrate on advanced issues such as finding feature selection algorithm in order to reduce large feature vectors obtained from the face and Iris fused vectors. This will actually reduce the time of execution of the system.

#### References

- [1] R. Connaughton, K.W. Bowyer, and P.J. Flynn. "Fusion of Face and Iris Biometrics from a Stand-Off Video Sensor". *MAICS*. April 2011.
- [2] J. Daugman. "How iris recognition works." Circuits and Systems for Video Technology, IEEE Transactions on 14.1 (2004): 21-30.
- [3] P. Viola, and J.J. Michael. "Robust real-time face detection." *International journal of computer vision* 57.2 (2004): 137-154.
- [4] M. Eskandari, T. Önsen, and D. Hasan. "A new approach for faceiris multimodal biometric recognition using score fusion." *International Journal of Pattern Recognition and Artificial Intelligence* 27.03 (2013): 1356004.
- [5] Y. Elmir, E.Zakaria, and A.Reda. "Score level fusion based multimodal biometric identification (Fingerprint & voice)." *Sciences of Electronics, Technologies of Information and Telecommunications (SETIT), 2012 6th International Conference on.* IEEE, 2012.
- [6] F. Wang, and J. Han. "Multimodal biometric authentication based on score level fusion using support vector machine." *Opto-electronics review* 17.1 (2009): 59-64.
- [7] Y. Wang, T. Tieniu, and A.K. Jain. "Combining face and iris biometrics for identity verification." Audio-and Video-Based Biometric Person Authentication. Springer Berlin Heidelberg, 2003.
- [8] Q. Tao and V. Raymond. "Threshold-optimized decision-level fusion and its application to biometrics." *Pattern Recognition* 42.5 (2009): 823-836.
- [9] W. Liu and L. Chong. "Face recognition based on rearranged modular 2DPCA." Advanced Intelligent Computing Theories and Applications. With Aspects of Artificial Intelligence. Springer Berlin Heidelberg, 2012. 395-403.
- [10] AT&T Laboratories Cambridge, The ORL Database of Faces. Internet: http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html. [2002].
- [11] CASIA. (2010). Chinese academy of sciences, CASIA Iris image database.internet: In <a href="http://biometrics.idealtest.org">http://biometrics.idealtest.org</a>, [2010]
- [12] G. Sutra, G.S. Sonia, and D. Bernadette. "The viterbi algorithm at different resolutions for enhanced iris segmentation." *Biometrics (ICB)*, 2012 5th IAPR International Conference on. IEEE, 2012.
- [13] M. Islam. "Feature and score fusion based multiple classifier selection for iris recognition." *Computational intelligence and neuroscience* 2014 (2014): 10.

# Design of a Virtual Simulator for Transesophageal Echocardiography Remote Control

José Zárate\* and Christina Pahl<sup>†</sup> \*<sup>†</sup>Biomechatronics Group, Ilmenau University of Technology Ilmenau 98693, Germany <sup>†</sup>Department of Clinical Sciences, Faculty of Biosciences and Medical Engineering, Universiti Teknologi Malaysia, Skudai 81310, Malaysia Email: \*jose-luis.zarate-moya@tu-ilmenau.de, <sup>†</sup>christina.pahl@tu-ilmenau.de

Abstract—Cardiovascular diseases are the leading cause of global mortality accounting for more than 17 million deaths annually. This represents approximately 30 % of all global mortality with an increasing trend and consequent growth in the use of cardiac imaging systems. Transesophageal echocardiography represents a highly convenient method to diagnose heart pathologies and to monitor cardiac procedures. However, the impact of radiation, musculoskeletal disorders and human error lead to the consideration of remote controlling this procedure. This work presents the virtual design and control of an endoscope system within a human esophagus. Results include the performance of linear and nonlinear 1 DOF system models. An approach for the future development of a novel robotic Transesophageal Echocardiography control is presented.

Keywords—Control, Endoscopic System, Virtual Simulator, Robotics, Transesophageal Echocardiography.

# I. INTRODUCTION

The first practical use (1976) of Transesophageal Echocardiography (TEE) in medical setups represented a novel acoustic window for cardiac procedures [1]. Due to the mortal effect of cardiovascular diseases, which accounts for 30 % of all global mortality with an increasing trend, TEE became a necessary cardiac imaging tool [2]. The TEE transducer is a modified endoscope with a set of piezoelectric crystals at its tip. By inserting the TEE probe into a human esophagus, cardiac structures can be visualized. The accuracy of the procedure is highly dependant on skills and experience to control the probe pose inside the esophagus under consideration of minimal contact with the surrounding tissue. These procedures become more complex if the cavity is compounded by close curves and irregular surfaces like the intestines during colonoscopy procedures [3]. In order to teach inexperienced operators to navigate this equipment, endoscopic simulators are used. There are four categories of this simulator type: mechanical simulators, animal models, cadaver models and computerized simulators [4].

# A. Problem Background

Radiation doses to patients and staff in the cardiac catheterization laboratory (cath lab) are relatively high when compared with conventional radiographic procedures as in radiology [5]. Furthermore, TEE often implies awkward hand positions for durations of 4 to 6 hours (coronary stenting or valve implantation). There are reports on unusually high incidence rates of musculoskeletal disorders [6]. Costs for appropriately trained TEE staff [7] are high due to the amount of training hours comprising at least 750 hours [8]. An approach towards the reduction of training hours [9] by supervising all TEE procedures based on mechatronic TEE control is of economic interest [10]. The employment of a remote navigation system removes the interventionists' dexterous and intuitive skills from the procedure.

# B. Related Works

This work refers to recently developed endoscopic remote control systems [11]. One approach represents the Multi-Angle Rear-Viewing Endoscopic Tool (MARVEL) [12]. It was developed for Minimally Invasive Neurosurgery (MIN) and consists of an articulated endoscope [13] equipped with a pivoting distal end [14]. This MIN tool is cable driven and allows to place the motors away from the joint leading to the reduction of surgical site occupation. Moreover, based on the same problem background of ensuring radiation-safe work environment, a remote catheter navigation system has been developed [15]. This model aims at providing a simple virtual simulator based on an endoscopic system. A Remote Catheter Navigation System (RCNS) has been developed to permit fluoroscopic X-ray guidance of percutaneous catheters from a radiation-safe location for the physician.

# C. Considerations

This paper presents the development of a TEE system using virtual simulation tools and control techniques. Several considerations have been taken into account. The esophagus is the chosen organ to be simulated, for its characteristics it can be simulated as a long cylindrical tube with regular shape. Depending on the height of the person, its approximate length is 25 cm [16]. Moreover, its approximate diameter is between 2 and 3 cm [17]. The model was built to be affected by two movements that displace the cylinder from its axis center. This effect is generated by the simulation of heartbeats, and the change of lung volume during respiration. For the representation of the TEE probe, a rigid articulated endoscope with a pivoting distal end was chosen in order to preserve the simplicity of the system. The endoscope meets resistance in the real model as it passes the upper esophageal sphincter resembled as the support point P of the endoscope, located in the beginning of the cylinder. P represents a start point from where the endoscope probe must be controllable until the tip using indicator coordinates without contact between shaft and esophageal surface under the exception of its distal end, which requires contact for the purpose of ultrasound wave transmission. Moreover, two position disturbances caused by heart and lung movements will be considered in order to avoid undesired contact of the probe shaft when passing esophageal structures.

### II. SYSTEM MODEL

In this section, the mechanical system model development, equations of movement that represent the endoscope displacement and complete system results defining in- and outputs for desired behaviour description will be explained.

#### A. Mechanical Model of the System

Based on the description provided in section I, the TEE probe is represented by an endoscope, which is designed to be introduced into the esophageal cavity with minimal surface contact. In order to introduce the tube, two axial forces are applied, which control the exact position and device navigation inside the cavity. However, it is possible using tracheal intubation at this stage. A second articulation is used to explore the target by using a pivot element with an imaging unit attached to the distal part. The representation of the system described previously is shown in Figure 1.

Further considerations have been taken into account to represent the system model as described below:

- The tube used in this system is considered to be rigid with a specific length *l* and mass *m*, where mass *S* is located at the center of the device.
- The vector  $\vec{P} = (\vec{i}, \vec{j}, \vec{k})$  represents the pivot point and resembles the tracheal intubation in form of a support point where the trajectory is adjusted to correct the navigation of the tube and corresponds with the system control.
- The orthogonal force components  $F_1$  and  $F_2$  provoke the correction of the position of the endoscope, represented by the generated angles  $\theta_1(t)$  and  $\theta_2(t)$ , respectively.
- The force  $F_3$  applied along the axis of the endoscope generates the displacement r(t) and corresponds to the navigation inside the cavity.
- The articulated element at the end of the endoscope has a length of t and its mass is negligible with respect to the complete device.

#### B. Equations of Movement and Dynamics

In order to obtain the dynamic model, the approach of equilibrium of forces set by the second Newton law, the equivalent with respect to rotation or Euler law can be utilized [18]. In this case, the mass is fixed frictionless in the center of gravity and the manipulated load is negligible. It should be noted that, apart of gravity and inertia forces, Coriolis and centripetal forces are involved due to the relative movement between the elements that compound the system and the configuration of the manipulator.



Fig. 1. Mechanical Model of the Endoscope.

Once the system parameters are set, the vector of the center of gravity  $r_l(t)$  is defined as shown in equation 1:

$$\vec{r_l}(t) = \begin{bmatrix} l_r(t)\cos\left(\theta_1(t)\right)\cos\left(\theta_2(t)\right) \\ l_r(t)\sin\left(\theta_1(t)\right)\cos\left(\theta_2(t)\right) \\ l_r(t)\sin\left(\theta_2(t)\right) \end{bmatrix}.$$
 (1)

where

$$l_r(t) = r(t) - \frac{l}{2} \tag{2}$$

Using Lagrange equations, based on energetic concepts, it is possible to obtain the relation of energies in the system referred to it in equation 3 [19]:

$$L = K - U \tag{3}$$

where  $K(\dot{r}_l(t))$  and  $U(r_l(t))$  are kinetic and potential energy, respectively. The result of both operations is shown in equations 4 and 5.

$$K(\dot{r}_{l}(t)) = \frac{1}{2}m\dot{r}_{l}(t)^{T}\dot{r}_{l}(t) = \frac{m}{16}\left(2(d-2r(t))^{2}\left(\dot{\theta}_{2}(t)^{2}+\dot{\theta}_{1}(t)^{2}\cos^{2}\left(\theta_{2}(t)\right)\right)+8\dot{r}(t)^{2}\right)$$
(4)

$$U(r_l(t)) = m\bar{g}r_l(t) = mg\left(r(t) - \frac{d}{2}\right)\sin\left(\theta_2(t)\right) \quad .$$
 (5)

The Lagrangian formulation is defined in equation 6, where  $q = [\theta_1(t) \ \theta_2(t) \ r(t)]^T$  are generalized system coordinates:

$$\frac{\mathrm{d}}{\mathrm{d}t} \left( \frac{\partial L}{\partial \dot{q}} \right) - \frac{\partial L}{\partial q} = Q. \tag{6}$$

This expression can be represented by the dynamic equation shown in 7, where M(q) is the inertial matrix,  $C(q, \dot{q})$  is the Coriolis and centripetal matrix and G(q) is the gravity matrix:

$$M(q)\ddot{q} + C(q,\dot{q}) + G(q) = Q.$$
 (7)

In this equation,  $Q = [\tau_1 \ \tau_2 \ F]^T$  represents the torque and force vector applied to the generalized coordinates q. However, the system is controlled by the forces applied to the handle of the endoscope. These forces are represented by the terms of  $F_1$ ,  $F_2$  and  $F_3$ , respectively. Moreover, they are related to the opposite end where the handle of the endoscope is located. At this point, it functions as a lever arm to control the navigation of the device inside the esophagus. This relation between both vectors is shown in expression 8:

$$Q = \begin{bmatrix} \tau_1 \\ \tau_2 \\ F \end{bmatrix}$$
$$= \begin{bmatrix} F_1(d - r(t))\cos(\theta_1(t))\cos(\theta_2(t)) \\ F_2(d - r(t))\cos(\theta_2(t)) \\ F_3 - F_1\sin(\theta_1(t))\cos(\theta_2(t)) - F_2\sin(\theta_2(t)) \end{bmatrix}$$
(8)

#### C. Nonlinear Model

Expression 7 is nonlinear. Therefore, it requires some calculations to obtain the direct dynamic model from it that provides the trajectory resulting from the application of a certain wrench Q. For this direct model as well as its subsequent use, any specific control technique may be utilized to obtain the dynamic model in state variables.

The natural state variables of the system are the positions and velocities of each joint, thus being the state vector  $x = [q \ \dot{q}]^T$ . Equation 7 can be expressed as:

$$\ddot{q} = -M(q)^{-1}(Q - C(q, \dot{q}) - G(q)).$$
(9)

By making use of the state vector, it is possible to obtain the dynamic model in state variables in equation 10:

$$\begin{bmatrix} \dot{q} \\ \ddot{q} \end{bmatrix} = f(q, \dot{q}) + g(q, \dot{q})u.$$
(10)

where

$$f(q, \dot{q}) = \begin{bmatrix} \dot{q} \\ M(q)^{-1}C(q, \dot{q}) + M(q)^{-1}G(q) \end{bmatrix}$$
(11)

$$g(q,\dot{q})u = \begin{bmatrix} 0\\ -M(q)^{-1}Q \end{bmatrix}$$
(12)

This dynamic model in state variables can be represented as a nonlinear space state system as shown in equation 13 [20]:

$$\dot{x} = f(x) + g(x)u \tag{13}$$

$$x = [q \ \dot{q}]^{\mathrm{T}} = \left[\theta_1(t) \ \theta_2(t) \ r(t) \ \dot{\theta_1}(t) \ \dot{\theta_2}(t) \ \dot{r}(t)\right]^{\mathrm{T}} \in \mathbb{R}^6$$
(14)

where  $u = [F_1 F_2 F_3]^T$  is the force input vector of the system.

#### III. CONTROL OF THE ENDOSCOPIC SYSTEM

## A. Linearization of the Nonlinear Model

The subsequent control design is simplified using linear methods. The position of the load as the output y of the system yields:

$$y = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = h(t) = \begin{bmatrix} r(t)\cos\left(\theta_1(t)\right)\cos\left(\theta_2(t)\right) \\ r(t)\sin\left(\theta_1(t)\right)\cos\left(\theta_2(t)\right) \\ r(t)\sin\left(\theta_2(t)\right) \end{bmatrix}.$$
(15)

Calculating the first derivative of the output results in y [20]:

$$\dot{y} = \mathcal{L}_f h(x) + \mathcal{L}_g h(x) u = \mathcal{L}_f h(x).$$
(16)

The acceleration of the output is obtained by differentiating once more:

$$\ddot{y} = \mathcal{L}_f^2 h(x) + \mathcal{L}_g \mathcal{L}_f h(x) u. \tag{17}$$

The system with the parametrization presented above is flat. The displacement of the system can be represented in cartesian coordinates by six states (14), therefore, the sum of the vectorial relative degree vrd has to be six as well. Hence, the values in this system are  $vrd = [2 \ 2 \ 2]$ . Using (15), it is possible to compute the measured load positions y by measuring state x, where it is possible to apply the same for load velocities  $\dot{y}$ , by (16). In this case, the second derivative  $\ddot{y}$  depends on the full input u. By the application of the general concept of feedback linearisation for MIMO systems, the expression is as follows [21],

$$y_j^{r_j} = \mathcal{L}_f^{r_j} h_j + \sum_{i=1}^m \mathcal{L}_{g_i} (\mathcal{L}_f^{r_j - 1} h_j) u_i .$$
(18)

The matrix  $\mathbf{M}(x)$ , defined as the decoupling matrix of the system, is given as:

$$\mathbf{M}(x) = \begin{bmatrix} \mathcal{L}_{g_1} \mathcal{L}_f^{r_1 - 1} h_1 & \cdots & \mathcal{L}_{g_m} \mathcal{L}_f^{r_1 - 1} h_1 \\ \vdots & \ddots & \vdots \\ \mathcal{L}_{g_1} \mathcal{L}_f^{r_m - 1} h_m & \cdots & \mathcal{L}_{g_m} \mathcal{L}_f^{r_m - 1} h_m \end{bmatrix}.$$
 (19)

The nonlinear system in (13) has a defined vector relative degree  $(r_1, r_2, \ldots, r_m)$  at the point  $x_0$  if  $\mathcal{L}_{g_i} \mathcal{L}_f^k h_i(x_0) = 0$  for  $0 \le k \le r_i - 2$  and  $i = 1, \ldots, m$  given that the matrix  $\mathbf{M}(x_0)$  is nonsingular. If the vector relative degree  $(r_1, r_2, \ldots, r_m)$  is well defined then (18) can be written as:

$$\begin{bmatrix} y_1^{r_1} \\ y_2^{r_2} \\ \vdots \\ y_m^{r_m} \end{bmatrix} = \begin{bmatrix} \mathcal{L}_f^{r_1} h_1 \\ \mathcal{L}_f^{r_2} h_2 \\ \vdots \\ \mathcal{L}_f^{r_m} h_m \end{bmatrix} + \mathbf{M}(x) \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_m \end{bmatrix}$$
(20)

The equation (20) is the preparation of a implicite linearization and since  $\mathbf{M}(x_0)$  is nonsingular, subsequently  $\mathbf{M}(x) \in \mathbb{R}^{m \times m}$  is nonsingular as well in a neighbourhood x of  $x_0$ . As a consequence, the control vector may be chosen as:

$$\begin{bmatrix} u_1\\u_2\\\vdots\\u_m \end{bmatrix} = -\mathbf{M}^{-1}(x) \begin{bmatrix} \mathcal{L}_f^{r_1}h_1\\\mathcal{L}_f^{r_2}h_2\\\vdots\\\mathcal{L}_f^{r_m}h_m \end{bmatrix} + \mathbf{M}^{-1}(x) \begin{bmatrix} y_1^{r_1}\\y_2^{r_2}\\\vdots\\y_m^{r_m} \end{bmatrix}.$$
 (21)

The equation 21 can be expressed as follows:

$$u = \rho(x) + \gamma(x)v \tag{22}$$

where

$$\gamma(x) = -\mathbf{M}^{-1}(x), \quad u = \begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_m \end{bmatrix},$$

the nonlinear MIMO system in (13) is linearised, yielding

$$v = \begin{bmatrix} y_1^{r_1} \\ y_2^{r_2} \\ \vdots \\ y_m^{r_m} \end{bmatrix}, \quad \rho(x) = -\mathbf{M}^{-1}(x) \begin{bmatrix} \mathcal{L}_f^{r_1} h_1 \\ \mathcal{L}_f^{r_2} h_2 \\ \vdots \\ \mathcal{L}_f^{r_m} h_m \end{bmatrix}.$$
$$\dot{x}_y = \mathbf{A}_y x_y + \mathbf{B}_y v \tag{23}$$

with

$$\mathbf{A}_{y} = \operatorname{diag}(\mathbf{A}_{y_{1}} \dots \mathbf{A}_{y_{m}}) \quad \mathbf{B}_{y} = \operatorname{diag}(\mathbf{B}_{y_{1}} \dots \mathbf{B}_{y_{m}}), \quad (24)$$

where each term individually is given by:

$$\mathbf{A}_{y_i} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 0 & 0 & 0 & \dots & 0 \end{bmatrix}, \quad \mathbf{B}_{y_i} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}.$$

By feedback control law and state transformation, the feedback linearisation is acquired, resulting in a linearised system in form of a chain of integrators [20].



Fig. 2. Model of the Nonlinear System in MATLAB SIMULINK.

# B. Linear Quadratic Regulator Control

An advantage of the Linear Quadratic Regulator (LQR) optimal control method over pole-placement is that LQR provides an automated way of defining an appropriate state-feedback controller [22].

For the optimal regulator problem, it shall be considered that,

$$\dot{e} = \mathbf{A}e + \mathbf{B}v \tag{25}$$

determines the matrix  $\mathbf{K}$  of the optimal controller

$$v = -\mathbf{K}e(t) \tag{26}$$

to minimize the performance index

$$J = \int_0^\infty (e^* \mathbf{Q} e + v^* \mathbf{R} v) dt \tag{27}$$

where  $\mathbf{Q}$  is a positive-definite (or positive-semidefinite) Hermitian or real symmetric matrix and  $\mathbf{R}$  is a positive-definite Hermitian or real symmetric matrix. It is important to note that the second term on the right-hand side of (27) accounts for the 'energy consumption' of control signals. The matrices  $\mathbf{Q}$  and  $\mathbf{R}$  determine the relative importance of the error and the 'consumption of the associated energy'. It is assumed that the control vector u(t) is unconstrained.

In MATLAB, the command that solves the continuoustime, linear, quadratic regulator problem and the associated algebraic Riccati equation is lqr(A, B, Q, R). This command calculates the optimal feedback gain matrix **K**. Therefore, the goal for finding **K** is to choose appropriate values of the matrices **Q** and **R**.

#### IV. SIMULATION AND RESULTS

All results of this work have been developed in MATLAB and its simulation tool SIMULINK. Figure 2 presents the shape of the simulation. Each part represents the process of linearisation of the nonlinear system, starting with the nonlinear plant, the state transformation (diffeomorphism) and the linearization block  $u = \varphi(x) + \gamma(x)v$ . All these parts build the Linearised Plant block  $\dot{x}_y = \mathbf{A}_y x_y + \mathbf{B}_y v$  which can be seen in Figure 3.

To begin with system tests, the dimensions of endoscope and imaging unit have been defined in Table I. It was assumed that both are cylindrical bodies. These characteristics have been chosen only for the purpose of experimentation and testing of the control system and can be readjusted later.

With respect to the results obtained in LQR control, the matrices  $A_y$ ,  $B_y$ , Q and R are defined as follows:



Fig. 3. Model of the Linear System in MATLAB SIMULINK.



Fig. 4. Simulation of the Endoscopic System in VRealm Builder.

TABLE I. DIMENSION OF THE COMPONENTS IN THE SIMULATION

Element	Length l (cm)	Diameter d (cm)	Mass m (g)
Endoscope (Main Body)	40	1	250
Imaging Unit	2.5	1	-
Esophagus	25	5	-

These values are applied for the closed loop of the linearised system. The result to apply the command lqr(A, B, Q, R) is presented subsequently:



Fig. 5. Model of the Nonlinear System in X axis with and without disturbance because of the movement of the lungs in MATLAB SIMULINK.



Fig. 6. Model of the displacement in Z axis of the Linear System with and without disturbance because of the movement of the lungs in MATLAB SIMULINK.

		$\mathbf{K_r}$	= [ K	r1	$\mathbf{K_{r2}}$ ] =	=	
Γ	100	0	0	17.32	0	0	1
	0	100	0	0	17.32	0	.
L	0	0	100	0	0	17.32	

With respect to the values obtained in the simulation, disturbances from heart an lungs were considered as a sinusoidal effect in z axis. Figure 5 shows a tendency to fix the position in the settled value. Moreover, Figure 6 presents the effect of disturbances in z axis, trying to follow the trajectory as close as possible. It is presenting a lag because of the response time of the driver due to the value of the disturbance frequency of 0.8 Hz. This value is higher than the typical respiratory rate of 12-18 breaths per minute in adults [23]. The effect of the disturbance due to the heartbeats is low and its effects will be studied in future research.

# V. DISCUSSION

Results show a simple TEE model being navigated under the impact of disturbances in a low complexity esophagus model. From Figure 5 it can be observed that the nonlinear system model shows a lower average error rate  $\Delta E(x) = 0.8$  mm than the linear one in Figure 6 with  $\Delta E(z) = 5.08$  mm. This is because the linear model represents a simplification of the more precise nonlinear approximation of the real model. The magnitude of the error in the nonlinear model is lower than the approximate image resolution in 3.5 MHz ultrasound image data with 0.9 mm. An adequate error magnitude in reference to ultrasound image resolution was the goal of this model. Moreover, the pivot point P, which is considered to be rigid in this paper needs to be modified to a flexible spring system. This is because P represents a flexible bundle of muscles in the real model and affects the forces applied in the current implementation in form of a damping system. However, this solution differs from related works mentioned in subsection I-B. Especially, the control approach and target parameters differ from the model presented in this paper. The design is similar to the MARVEL system due to the simplicity of endoscopic system geometry. Individual adaptions towards control [24] have been described in previous works [25]. Moreover, incorporating individual patient data using electronic health records [26] enables best possible trajectory planning [27].

#### VI. CONCLUSION

Based on its full actuated first order dynamics, the 1 DOF endoscope system has been modelled. Its inputs represent forces applied at the handle of the endoscope. The generated trajectory control was developed using LQR. It shows the behaviour of the system affected by external disturbances, which are heartbeat and respiration. With respect to external perturbations, sinusoidal signal was used to simulate the effect of the heart and lungs in the esophagus. Moreover, results show that LQR has shown sufficient performance, especially, in driving the accelerations in the system to adjust its displacement. Subsequent research encloses the development of a robotically performed TEE with higher DOF.

#### ACKNOWLEDGEMENT

The authors would like to gratefully acknowledge financial support from Technische Universität Ilmenau.

#### REFERENCES

- [1] J. B. Seward, B. K. KHANDHERIA, J. K. OH, M. D. ABEL, R. W. HUGHES, W. D. EDWARDS, B. A. NICHOLS, W. K. FREEMAN, and A. J. TAJIK, "Transesophageal echocardiography: technique, anatomic correlations, implementation, and clinical applications," in *Mayo Clinic Proceedings*, vol. 63, no. 7. Elsevier, 1988, pp. 649–680.
- [2] A. Alwan et al., Global status report on noncommunicable diseases 2010. World Health Organization, 2011.
- [3] A. Haycock, A. D. Koch, P. Familiari, F. van Delft, E. Dekker, L. Petruzziello, J. Haringsma, and S. Thomas-Gibson, "Training and transfer of colonoscopy skills: a multinational, randomized, blinded, controlled trial of simulator versus bedside training," *Gastrointestinal endoscopy*, vol. 71, no. 2, pp. 298–307, 2010.
- [4] D. J. Desilets, S. Banerjee, B. A. Barth, V. Kaul, S. R. Kethu, M. C. Pedrosa, P. R. Pfau, J. L. Tokar, S. Varadarajulu, A. Wang *et al.*, "Endoscopic simulators," *Gastrointestinal Endoscopy*, vol. 5, no. 73, pp. 861–867, 2011.
- [5] O. Morrish and K. Goldstone, "An investigation into patient and staff doses from x-ray angiography during coronary interventional procedures," *The British journal of radiology*, 2014.
- [6] H. E. Vanderpool, E. A. Friis, B. S. Smith, and K. L. Harms, "Prevalence of carpal tunnel syndrome and other work-related musculoskeletal problems in cardiac sonographers." *Journal of occupational and environmental medicine*, vol. 35, no. 6, pp. 604–610, 1993.

- [7] Y. N. Aditya, H. N. Abduljabbar, C. Pahl, L. K. Wee, and E. Supriyanto, "Fetal weight and gender estimation using computer based ultrasound images analysis," *International Journal of Computers*, p. 11, 2013.
- [8] M. K. Cahalan, M. Abel, M. Goldman, A. Pearlman, P. Sears-Rogan, I. Russell, J. Shanewise, W. Stewart, and C. Troianos, "American society of echocardiography and society of cardiovascular anesthesiologists task force guidelines for training in perioperative echocardiography," *Anesthesia & Analgesia*, vol. 94, no. 6, pp. 1384–1388, 2002.
- [9] C. Pahl, E. Supriyanto, N. H. B. Mahmood, and J. Yunus, "Cervix detection using squared error subtraction," in *Modelling Symposium* (AMS), 2012 Sixth Asia. IEEE, 2012, pp. 121–125.
- [10] C. PAHL, "Novel method for autonomous ultrasound cervix scanning," Ph.D. dissertation, UNIVERSITI TEKNOLOGI MALAYSIA, 2012.
- [11] D. C. Kho, C. Pahl, E. Supriyanto, Y. M. Myint, A. A. Faudzi, and M. I. Salim, "Motorized remote control of transesophageal echocardiography (tee) probe tip: Preliminary testing," *View in Article*, 2014.
- [12] M. Shearn, S. Y. Bae, R. Korniski, H. Manohara, J. Mondry, and H. Shahinian, "Multi-angle rear-viewing endoscopic tool (marvel) for minimally invasive neurosurgeries," *Journal of Medical Devices*, vol. 6, no. 1, p. 017540, 2012.
- [13] C. Castro, S. Smith, A. Alqassis, T. Ketterl, Y. Sun, S. Ross, A. Rosemurgy, P. P. Savage, R. D. Gitlin *et al.*, "Marvel: A wireless miniature anchored robotic videoscope for expedited laparoscopy," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 2926–2931.
- [14] Y. Bae, A. Liao, H. Manohara, and H. Shahinian, "Adjustable-viewingangle endoscopic tool for skull base and brain surgery," 2008.
- [15] Y. Thakur, C. J. Norley, D. W. Holdsworth, and M. Drangova, "Remote vs. manual catheter navigation: a comparison study of operator performance using a 2d multi-path phantom," in *SPIE Medical Imaging*. International Society for Optics and Photonics, 2009, pp. 72611A– 72611A.
- [16] D. C. Allen, R. I. Cameron, and M. B. Loughrey, "Esophagus," in *Histopathology Specimens*. Springer, 2013, pp. 13–22.
- [17] L. J. Skandalakis and J. E. Skandalakis, Surgical anatomy and technique: a pocket manual. Springer Science & Business Media, 2013.
- [18] J. L. Z. Moya, "Diseño y control de un sistema móvil con brazo robot para aplicaciones de manipulación de objetos peligrosos basado en la teleoperación," Ph.D. dissertation, 2005.
- [19] J. J. Craig, Introduction to robotics: mechanics and control. Pearson Prentice Hall Upper Saddle River, 2005, vol. 3.
- [20] A. Isidori, *Nonlinear Control Systems*, 3rd ed., M. Thoma, E. D. Sontag, B. W. Dickinson, A. Fettweis, J. L. Massey, and J. W. Modestino, Eds. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1995.
- [21] K. L. Knierim, K. Krieger, and O. Sawodny, "Flatness based control of a 3-dof overhead crane with velocity controlled drives," in *Mechatronic Systems*, 2010, pp. 363–368.
- [22] K. Ogata, *Modern control engineering*, 5th ed., ser. Prentice-Hall electrical engineering series. Instrumentation and controls series. Boston: Prentice-Hall, 2010.
- [23] W. F. Ganong and K. E. Barrett, *Review of medical physiology*. McGraw-Hill Medical ^ eNew York New York, 2005, vol. 21.
- [24] C. Pahl and E. Supriyanto, "Guide to fuzzy logic based transesophageal echocardiography motorized probe control," in *Methods and Models in Automation and Robotics (MMAR), 2015 20th International Conference* on. IEEE, 2015, pp. 892–896.
- [25] S. E. Pahl, Christina, "Personalized cervix ultrasound scan based on robotic arm," *PSRC-Planetary Scientific Research Center Proceeding*, 2012.
- [26] C. Pahl, M. Zare, A. B. Ahmad, V. Detschew, D. Ammon, S. Lehnert, and E. Supriyanto, "Identification of quality parameters for an e-health platform in the federal state of thuringia in germany," *Journal of Soft Computing and Decision Support Systems*, vol. 1, no. 1, pp. 17–23, 2014.
- [27] C. Pahl, M. Zare, M. Nilashi, M. A. de Faria Borges, D. Weingaertner, V. Detschew, E. Supriyanto, and O. Ibrahim, "Role of openehr as an open source solution for the regional modelling of patient data in obstetrics," *Journal of biomedical informatics*, vol. 55, pp. 174–187, 2015.

# Emotional Feedback for Service Robots using Hapto-Acoustic Interface

Christina Pahl Biomechatronics Group Ilmenau University of Technology 98693 Ilmenau, Germany christina.pahl@tu-ilmenau.de

behavior-oriented human-machine Abstract—A communication system based on bio-inspired information processing is developed in this paper. A low cost acoustic interface using microphones is created as an emotional interaction channel for robotic assistants in domestic environments for the care of elderly people with mild cognitive impairment. Classification of communication signals such as striking, petting, speech and noise could be successfully realized using 3 condenser microphones and deployed on a robot prototype. The proposed framework comprised of an electronic microphone interface circuit, an ATmega128 microcontroller for processing of signals from the circuit, a Fast Fourier Transform (FFT) algorithm implementation on the microcontroller for Fourier analysis, LUFA library for USB data transfer from the microcontroller, a machine learning module based on Support Vector Machine (SVM) for signal classification and an amplifier based on an RFT transistor capsule for feedback. Results show that in all four cases, a successful differentiation with reliable and robust deviation between the signals can be made. The success rate with the SVM classifier shows with 87.5 % a highly promising classification result. It can be concluded that this novel approach is highly useful for non-verbal human-machine communication.

Keywords—acoustics; communication; emotional feedback; elderly care; haptics; human-machine interaction; service robots.

# I. INTRODUCTION

### A. Background

Recent results have shown that robots can serve as efficient companions for elderly people with Mild Cognitive Impairment (MCI) living alone at home [1]. The spectrum of required assistive functionalities of such a robot companion is broad. It spans applications such as reminding functions e.g. taking medication, drinking or scanning specific body organs [2]. Moreover, cognitive stimulation exercises via mobile videophony with relatives or caregivers is important for the detection and evaluation of critical situations, like falls 3]. This paper addresses the requirements of human-robot communication on such robots [4]. Previous works in this direction include vision based feedback, verbal feedback to action executed, as well as responses based on touch [5]. For example, Aldebaran Nao robots offer touch sensors on head, hands and feet that can be used to provide feedback [6].

# Karthik Mahesh Varadarajan

Automation and Control Institute Technical University of Vienna 1200 Vienna, Austria kv@acin.tuwien.ac.at

Nao has been used for the treatment of autism kids in domestic environments as part of the ASK program. But, typical feedback actions on such haptic systems are restricted to tapping or patting [7]. Gesture detection based on visual feedback has also been used in a multitude of systems geared towards rehabilitation robotics [8]. The most significant of rehabilitation robots include MIT-Manus, Bi-Manu Track and MIME that provide assistive exercising. Exoskeletons such as HULC, HAL, eLEGS help augment dexterous manipulation and locomotion in physically challenged users. Moreover, studies have been carried out on the interpretation of social touch on an artificial arm [9]. The modality of touch was correctly classified in approximately 71 % of the trials using EIT-based sensitive skin. This shows that classification of social touch can be achieved with accuracies comparable to those achieved by humans [10].

# B. Social Robotics and Elderly Care

In the context of communication and aid of elderly people, machines are being developed that are specialized in interacting socially [11]. This generation of robots is known as social robots and are characterized by their ability to interact with people according to their needs [12]. Specifically, it has been documented that a higher proportion of elderly people suffer from loneliness and its consequences such as marginalization, dementia and depression [13]. Inadequate interpersonal relationships increase these effects [14].



Fig. 1. SCITOS G5 Senior citizen assistant robot developed by TUIL. (Source: TUIL)

A successful therapeutic approach, the interpersonal psychotherapy (IPT) was developed [15]. Robots, such as Kaspar or PARO are mechatronic therapists, which are used in clinical environments and homes for targeted groups of elderly citizens. With the aid of robotic assistance systems, aged people can gain back their independence. In Fig. 1, two aged subjects are shown interacting with the SCITOS G5 Senior citizen assistant social robot. The general acceptance of robots is small, but rapidly increasing [16]. Due to demographic changes in developed countries, the market for social robots is rapidly growing with greater deployment of social robots in regular and old-age homes.

# C. Haptic Sensors for Social Robots

Table I compares four examples of social robots: Kaspar, the Haptic Creature, Paro and RI-MAN. Kaspar is geared at therapies for autistic patients, whereas Paro supports treatments for patients suffering from Dementia, Alzheimer's and other cognitive disabilities. Kaspar, which is utilized for treating autistic children, employs capacitive haptic sensors. These capacitive sensors are flexible and integrated in a software that has predefined action commands for different touch patterns. Paro's appearance resembles a baby seal. The robot responds to petting actions with emotions. The Haptic Creature has 56 Force-Sensing Resistors (FSR) integrated in its body. This robot realizes its haptic interactions using dielectric material between the electrodes. Alternately, Quantum Tunneling Composites (QSRs) have also been used. Finally, RI-MAN is a service robot targeted for the movement of hospitalized patients. RI-MAN has flexible sensors which are attached on polymers. These haptic robotic systems utilize electronic sensors for touch sensing. While modeling of typical haptic actions of average humans that include the tap, stroke, scratch, light touch, tickle have been studied through social robots, such as the Kaspar, the Haptic Creature, Paro and the RI-MAN among others, there is a definitive lack of research aimed at modeling emotional feedback from the surface of social robots for the purpose of emotional interaction with elderly users. Capacitive artificial skin such as in the case of Kaspar, is one approach towards haptic communication. Capacitive artificial

skin, which forms the current state of the art in social robots, shows that these materials tend to be more flexible and light weight. Modifiable material properties enable complex motor controls. Tactile Sensors are essential for robots which are interacting with the environment. Tactile interfaces protect against damages of the robot and also its environment. Currently, research challenges in this area lie in the development of flat and flexible sensor technology. This includes the development of capacitive sensors in form of small brushes as a tactile sensor technology for humanoid robots. Such capacitive pressure sensors find application in humanoids like iCub [17]. Nevertheless, such sensors are expensive and not easily scalable or modifiable. These sensors also require rigid grid configurations and calibration for the detection of pre-learnt emotional haptic feedback response patterns. Alternatively, it is possible to design haptic sensors based on magnetic or acoustic sensors that are relatively cheap. Magnetoelastic sensors consist of two mutually orthogonal coils in a ferromagnetic material. Out of this literature review, it can be concluded that there is a growing demand for haptic interfaces in socially interacting robots. However, there is no hapto-acoustic approach documented which is based on microphones.

# II. PROPOSED APPROACH

This paper addresses limitations and concerns with traditional haptic sensing technologies by proposing a novel, inexpensive, scalable haptic sensor. Various emotional haptic feedback responses that are suited to use for elderly respondents (striking, petting, speech) are modeled using the designed haptic sensor. The test environment at this stage was designed to be without any ego-motion noise (motors) causing vibrations. The position of the tactile input was given by a marker above the sensor arrangement defined by the triangulation of sensors. The first 64 values of FFT were utilized because of their high information content which makes them relevant for system representation. Training data for the SVM were collected under ideal conditions with repetitive input patterns for each data class. The amount of training data was defined by 10 signal inputs per class using a linear SVM.

Robotic System	Kaspar	Haptic Creature	Paro	RI-MAN
Application	Autism Therapy	Therapy, Entertainment	Therapy for People with Dementia	Healthcare, Movement of hospitalized patients
Implementation	Flexible Sensors attached on polymers	56 FSR on Fiber-glass	Dielectric Material between three Electrodes	Flexible Sensors attached on polymers
Realization	Capacitive Sensors	Capacitive Sensors	Patent: PCT/JP02/05227	Sensor system with Semi-conductor

 TABLE I.
 HAPTIC SENSORS IN SOCIAL ROBOTS

The paper also provides a framework that can be expanded to a range of haptic gestures that can be used as feedback to companion robots. Specifically, a novel acoustic gesture detection system is investigated as a possible modality of human-robot communication in home environments. The emotional HRI based on acoustic sensing designed in this paper is extremely inexpensive compared to other traditional haptic sensors. Although, this work uses commonly available components (microphones), the application by itself represents an innovative approach. The general framework for such an acoustic haptic sensor is presented in this paper along with details of implementation and corresponding results.

The system includes a prototype of a robot head, a microphone amplifier circuit and a signal conditioning processor. The setup was tested at the Department of Neuroinformatics and Computer Science at the Ilmenau University of Technology (TUIL) using the SCITOS G5 robot from MetraLabs, running on MIRA middleware. The proposed communication interface based on haptic interface was attached to the robot prototype externally. This contains a robust acoustic sensor which was integrated into the automated system and tested successfully. Results from haptic stimulus processing on the system are presented and discussed. The evaluation also addresses demerits of the proposed system and possible future enhancements to the approach. Acoustic sensors can be used to register minimal forces for tactile communication. The robot receives any information in form of vibrations, which results in fine sound waves. Microphones attached to the inner surface material can receive this energy. There are different types of microphones. These include laryngeal microphone, piezo or crystal microphone, electret condenser microphone and moving coil microphone among others. Out of these, the condenser microphone is the most suitable for this setup, since it is reasonable in terms of costs and offers a compact design. Furthermore, it is robust in its operation compared to other microphones. In Fig. 2 the top view on the arrangement of the acoustic sensors can be obtained. The triangle shows the frontal part of the spherical robot prototype head.



Fig. 2. Sensor arrangement on robotic prototype head.



Fig. 3. Sensor arrangement on robotic prototype head.

This setup consists of 3 condenser microphones with a diameter of 9.7 mm in an arrangement of 120 degrees distance to each other. Recording of the acoustic output was carried out using a PicoScope 3000 oscilloscope and logic analyzer.

# A. Haptic Information Patterns

Due to the nature of tactile information, it is important to understand these information in terms of the following parameters: vibration, pressure, temperature and type of touch. Exploring these parameters (while being an ideal approach in terms of performance) would be very complex since many different types of sensors would be required. In order to meet the objective to design an emotional feedback circuit for the classification of different tactile data, it suffices to recognize haptic information only in this research. Since petting of the robot head represents a shear force applied on the surface of the material, it can theoretically be distinguished easily from striking force resulting in different vibration characteristics. Striking represents a negative feedback and can be seen as a short and strong pressure on the robot head surface. Fig. 3 represents striking in form of 4 red vertices without any change of position with advance of time. The grid represents a reference for pattern motion. It underlines the difference in touch characteristics. Petting is represented here as the change of contact position with time. The green vertices mark the positions during two different moments,  $t_1$  and  $t_2$ . The three condenser microphones in the emotional acoustic feedback measurement circuit, placed at strategic positions on the robot head pick up the signal from the haptic actions before being amplified and fed into the signal processing circuitry. The amplifier circuit comprising the RFT transistor capsule that is used between the microphones and the acoustic signal processing microcontroller is shown in Fig. 4 (along with its hardware realization and component identification). Different prototypes representing the robot head surface were used in experiments for the ideal placement of acoustic sensors and material.



Fig. 4. Sensor arrangement on robotic prototype head.

#### B. Acoustic Signal Processing

The acoustic signals from the microphones are then processed on a dedicated microcontroller. A dedicated ATmega128 microcontroller was chosen because it offers sufficient computational power to fulfil the requirements of this framework. The signal processing hardware chain inside the microcontroller provides an ADC for the digitalization of the analog input from the electronic circuit. The ADC quantizes the input signals with a resolution of 10 bits. Since the maximum frequency of usable input signal is 200 kHz, this range is sufficient for reliable representation of the audio signals. After discretization of the signal input, a 128 point Fast Fourier Transform (FFT) is applied on the signal by the microcontroller. This eliminates the DC component. 64 of the 128 values are significant, while the rest are insignificant. This is as a result of restricting the input signal to the real domain. As pre-processing for the FFT-algorithm, the values of the signal input are arranged so that signal values lie on even indices of the array while array locations with odd indices are filled with zeros. The FFT-algorithm then recognizes all even indices as real numbers and all odd indices as imaginary units The Danielson-Lancoz-Algorithm that of the signal. decomposes the 128 point DFT into a cascade of several 2 point DFTs is then applied, in order to simplify the computation of the FFT. The estimated Fourier coefficients are sent via USB using the LUFA microcontroller library to the host-system G5 running MIRA.

# C. Classification

Data processing of tactile information is traditionally based on cognitive maps. For the classification of tactile information some groups like the Kaspar autism treatment research group use histogram based classification. The classification of spectral data was performed on the host system MIRA. This system offers neural networks and SVM, among other libraries for classification. MIRA offers different "units" or modules that represent specific classes. The communication between

these units is done using channels. One unit may write data, which is then read by other units that are related to the channel. For this project, two units were used. One unit, being the driver unit and the other unit for classification. A Support Vector Machine (SVM) was used to distinguish between the different haptic Fourier coefficient streams. The SVM received 64 inputs which were classified on the basis of training data vectors. However, this module is directly provided by MIRA and only required structural adaptations. The driver unit was used to create a connection between the microcontroller and MIRA using LibUSB. Data transferred by the USB is then written on a channel, which the classification unit accesses. During the training phase, the driver unit writes data for training purposes of the SVM. The classification unit waits until new coefficients are written on the channel before a categorization can be done in combination with the trained SVM. The coefficients are then labeled to one of the given categories (petting, striking, speech or noise). For the purpose of striking classification, an amplitude controlled trigger was used in combination with the microcontroller so that in this case only striking spectra passed the channels whereas data from other signal components were filtered. This represents a preprocessing step in order to facilitate the classification step. After training the SVM with training data, a validation data set was used to define the correctness of the SVM. Therefore, the number of correctly classified vectors were compared with the complete number of vectors of the data set. This result was calculated in all data classes used. In order to avoid strong variations in the spectra pattern for all classes and subsequent generation of weak support vectors, averaging of consecutive spectra has been applied. This method stabilized the pattern and minimalized the coefficients of variance. This analysis showed that averaging over 4 values significantly reduced variations among the coefficients compared to the utilization of single spectra. In order to realize this step, the driver class had to be extended. This lead to the generation of a new channel value within the USB interface for subsequent input values. The trigger of the microcontroller therefore required to send 4 subsequent spectra after striking. Box I shows the typical signal component obtained during training (the maximal amplitude component and frequency with highest Fourier coefficient) are listed. As mentioned earlier, classification of these component values was carried out using SVM.

#### D. Classification

After processing acoustic sensor signals, a reaction from the robot in the form of a feedback was performed to convey the successful acknowledgement of the feedback to the user. This also initiates further interaction with the user. Human interpretation of emotional feedback characteristics was used in this situation.

(a): $A_{max} = 1.1V$ , $f = 14 \text{ mHz}$
(b): $A_{max} = 2.2V$ , $f = 9 \text{ mHz}$
(c): $A_{max} = 0.3V$ , $f = 56 \text{ mHz}$
(d): $A_{max} = 0.1V$ , $f = 42 \text{ mHz}$

Box. I. Typical signal component (highest amplitude and frequency of highest Fourier component) obtained during training

The usage of actuation as a robotic reaction in response to a human action is a feasible approach, but this requires an actuator. Moreover, using acoustics can be a more viable approach. However, depending on the situation, an acoustic response is sometimes not desired and therefore does not represent the best possible solution. To circumvent the situation, it should be possible to mimic a more human-like reaction to the presented feedback. In this setup, the control of the robot's eyes represents the best solution. Here, the emotion model is based on two conditions - either positive or negative feedback. In the condition resulting in a reward, realized by a petting gesture on the surface of the robot, the eyes of the robot are programmed to react towards the haptic information, followed by a verbal acknowledgement. A similar, but discriminative response to a negative feedback can be modeled as well. Moreover, incorporating individual user data within electronic health records on the same robotic platform enables user context interaction planning [18].

# III. RESULTS

All three microphones were connected to the amplifier circuit, microcontroller and processed by the classification algorithm. In order to understand the signal behavior of the framework, four different acoustic situations were tested. These signals corresponded to silence, speech, gentle touch in form of a single petting gesture and striking.



(a) Output of electronic circuit for striking



(c) Output of electronic circuit for petting



TABLE II. CONTINGENCY TABLE FOR TOUCH CLASSIFICATION USING 40 DATA INPUTS

		Predicted				
_		Petting	Striking	Speech	Noise	
	Petting	9	1	0	0	
s	Striking	1	8	0	1	
ctu las	Speech	0	0	9	1	
A O	Noise	0	0	1	9	

Table II displays the performance of this classification method using SVM. Each column of the matrix represents the instances in a predicted class, while each row represents the instances in an actual class. It can be observed that 87.5 % correct classifications of 10 classifications trials for each of all four cases have been achieved.

#### A. Evaluation

In all four cases, a high differentiation could be observed. Here, the values refer to the first 10 *ms* of signal response. In the first case, a stroke evokes a signal with high amplitude characteristics and low frequency. During the first 5 *ms* the signal shows low frequency and very high amplitude, while after 10 *ms* the characteristics switch to low amplitude and high frequency with unstable response. In the case of speaking, there is a strong periodic and continuous wave with high amplitude and very low frequencies. Silence has continuous signal characteristics with low amplitudes.



(b) Output of electronic circuit for speech



(d) Output of electronic circuit for noise

Highly differentiable characteristics can be seen here. Less differentiable are silence and petting. Here, the differentiation can be made on the basis of amplitude. The analog outputs of the acoustic sensors, corresponding to the various modalities of striking, petting, speech and noise are shown in Fig. 5. For a total of 40 user samples, classification showed a differentiation value of almost 90 % according to the following distribution:

Petting:	0.9
Striking:	0.8
Speech:	0.9
Noise <sup>.</sup>	09

Petting has been correctly assigned in 9 out of 10 cases, which also holds true for Speech and Noise. Striking, however, was wrongly classified in 2 out of ten cases. It can be observed that the usage of the acoustic sensors is highly suitable for human-machine interaction.

# IV. DISCUSSION

These results show with almost 90 % classification correctness, significantly higher classification rates than some alternative approaches but still in the same range as other robust online touch pattern recognition solutions with an accuracy for hit being as high as 87.83 % [20]. The here achieved signal resolution is sufficient enough to perform binary classifications for all input data [10]. Merits of this approach are the low-cost, simple realization and training due to the strong characteristics of signal inputs. Disadvantages lie in possible sensitivity issues of the amplifier circuit due to unclean solder connections. Moreover, the choice of reference voltage may lead to faulty signal processing with subsequent low signal resolution. Furthermore, robustness may get improved by a different choice of data package and data type as well as the use of further information like temperature of touch. This approach, and its interface might be deployed in human-robot interaction where simple gestures replace complicated communication context recognition and provide a channel to express emotions by touch. This approach might get employed for adaptive behavior control of HRI.

#### V. CONCLUSION

In this paper, acoustic sensors were utilized for detecting affective haptic inputs representing a novel approach. This includes a system with higher classification rates compared to previous work to process hapto-acoustic data. The proposed method can be used for the recognition of typical feedback from elderly subjects and is thus well suited as a low-cost, scalable, alternative towards haptic feedback systems for domestic and elderly care robots. Future work will involve further enhancing the gamut of gestures that can be recognized and differentiated by the system.

#### ACKNOWLEDGMENT

The authors would like to gratefully acknowledge the support of Technische Universität Ilmenau for providing financial assistance.

#### REFERENCES

- H. Gross, C. Schroeter, S. Mueller, M. Volkhardt, E. Einhorn, A. Bley, T. Langner, "Further progress towards a home robot companion for people with mild cognitive impairment." In Systems, Man, and Cybernetics (SMC), IEEE, 2012, pp. 637-644.
- [2] C. Pahl, E. Supriyanto, "Personalized Cervix Ultrasound Scan Based On Robotic Arm." In International Conference on Systems and Electronic Engineering (ICSEE'2012), 2012, pp. 18-19.
- [3] C. Granata, et al. "Voice and graphical-based interfaces for interaction with a robot dedicated to elderly and people with cognitive disorders." RO-MAN, IEEE, 2010, pp. 785-790.
- [4] S. M. Rabbitt, A. E. Kazdin, and B. Scassellati, "Integrating Socially Assistive Robotics into Mental Healthcare Interventions: Applications and Recommendations for Expanded Use." Clinical Psychology Review, 2014, pp. 35-46.
- [5] G. van den Broek, F. Cavallo, and C. Wehrmann, eds. AALIANCE ambient assisted living roadmap. Vol. 6. IOS press, 2010.
- [6] A. Csapo, E, Gilmartin, J. Grizou, J.G. Han, R. Meena, D. Anastasiou, K. Jokinen, and . Wilcock. "Multimodal conversational interaction with a humanoid robot." In Cognitive Infocommunications (CogInfoCom), 2012 IEEE 3rd International Conference on, 2012, pp. 667-672.
- [7] J.J. Cabibihan, I. Ahmed, and S. Sam Ge. "Force and motion analyses of the human patting gesture for robotic social touching." In Cybernetics and Intelligent Systems (CIS), 2011 IEEE 5th International Conference on, IEEE, 2011, pp. 165-169.
- [8] A. Bulling, U. Blanke, and B. Schiele. "A tutorial on human activity recognition using body-worn inertial sensors." ACM Computing Surveys (CSUR) 46, no. 3, 33, 2014.
- [9] D. S. Tawil, D. Rye, and M. Velonaki. "Interpretation of the modality of touch on an artificial arm covered with an EIT-based sensitive skin." The International Journal of Robotics Research, 31, no. 13, 2012, pp. 1627-1641.
- [10] D. S.Tawil, D. Rye, and M. Velonaki. "Interpretation of social touch on an artificial arm covered with an EIT-based sensitive skin." International Journal of Social Robotics 6, no. 4, pp. 489-505, 2014.
- [11] M. Díaz, J. Saez-Pons, M. Heerink, and C. Angulo, "Emotional factors in robot-based assistive services for elderly at home." In RO-MAN, 2013 IEEE, 2013, pp. 711-716.
- [12] R. Bemelmans, G. J. Gelderblom, P. Jonker, and L. De Witte, "Socially assistive robots in elderly care: A systematic review into effects and effectiveness." Journal of the American Medical Directors Association 13, no. 2, 2012, pp.114-120.
- [13] E. Mordoch, A. Osterreicher, L. Guse, K. Roger, and G. Thompson, "Use of social commitment robots in the care of elderly people with dementia: A literature review." Maturitas 74, no. 1, 2013, pp. 14-20.
- [14] R. Khosla, M. T. Chu, and K. Nguyen, "Enhancing Emotional Well Being of Elderly Using Assistive Social Robots in Australia." In Biometrics and Kansei Engineering (ICBAKE), 2013 International Conference on, IEEE, 2013, pp. 41-46.
- [15] S. M. Gunlicks, and M. M. Weissman, "Interpersonal Psychotherapy (IPT)." Handbook of Interpersonal Psychology: Theory, Research, Assessment, and Therapeutic Interventions, 2011, pp. 533-544.
- [16] C. Pahl, E. Supriyanto, N.H.B. Mahmood, Y. Yunus, "Cervix detection using squared error subtraction." In Modelling Symposium (AMS), 2012 Sixth Asia, IEEE, 2012, pp. 121-125.
- [17] A. Parmiggiani, Alberto, M. Maggiali, L. Natale, F. Nori, A.Schmitz, N. Tsagarakis, J. S. Victor, F. Becchi, G. Sandini, and G. Metta. "The design of the iCub humanoid robot." International journal of humanoid robotics 9, no. 04, 1250027, 2012.
- [18] C. Pahl, M. Zare, M. Nilashi, M., M. A. de Faria Borges, D. Weingaertner, V. Detschew, V. & O. Ibrahim, "Role of OpenEHR as an open source solution for the regional modelling of patient data in obstetrics." Journal of biomedical informatics, 55, 2015, pp. 174-187.
- [19] Y. M. Kim, S. Y. Koo, J. G. Lim, D. S. Kwon, "A robust online touch pattern recognition for dynamic human-robot interaction." Consumer Electronics, IEEE Transactions on, 56 (3), 2010, pp. 1979-1987.

2015 Pattern Recognition Association of South Africa and Robotics and Mechatronics International Conference (PRASA-RobMech) Port Elizabeth, South Africa, November 26-27, 2015

# Control Considerations of the Low Cost Prosthetic Touch Hand

Drew van der Riet<sup>1</sup>, Riaan Stopforth<sup>2</sup>, Glen Bright<sup>3</sup> School of Mechanical Engineering, University of KwaZulu-Natal, Howard College, Durban, South Africa e-mail: <sup>1</sup>213569879@stu.ukzn.ac.za, <sup>2</sup>stopforth@ukzn.ac.za, <sup>3</sup>brightg@ukzn.ac.za

*Abstract*— The control of the Touch Hand: a low-cost electrically powered prosthetic hand is explored in this paper. A novel Haptic User Interface is proposed and tested as a supplement for amputee-prosthetic EMG control. The performance of the prosthetic hand is tested through gripping tests. The hand is capable of performing a lateral grip of 3.7 N, a power grip of 19.5 N and to passively hold a weight of up to 8 kg with a hook grip. The hand is also tested on an amputee and used to perform basic tasks. The amputee took 30 min to learn how to operate the hands basic gripping functions.

# *Keywords—electromyography; EMG control; haptic user interface; upper limb prosthetics; low-cost prosthetics*

#### I. INTRODUCTION

There are three main functions of the human hand: gripping, manipulating and exploring [1]. The hand also has to be aesthetical appealing, to have normal human appearance. A prosthesis is used to supplement these functions for those without hands. Currently available prostheses are a nonfunctional, aesthetical hand and functional prosthetic hands. These functional hands allow for varying degrees of gripping and manipulation ability. Current available prostheses, for manipulation purposes, range from non-functional aesthetic options to fully integrated prosthetic devices, which does not provide any form of sensory restoration. There has been significant progress towards sensory restoration in prosthetics which has been shown to aid in the control of prosthetics [2].

There are further active areas of study that include control of the prosthesis. The amputee's control of the prosthesis is one of the most important design criteria in prosthetics. Without adequate control options for the amputee, the advanced functionality of a prosthetic hand is wasted. Mechanical prosthetic hands use cables and leavers attached to the amputee's body to drive open and close the prosthetic hand. Electrical hands are controlled by the amputees mind. Signals can be taken invasively with the use of electrodes that are surgically placed in the nerves of the amputee's arm. Noninvasively approaches use an electroencephalography (EEG) headset to read brain signals. Another method to control the prosthetic hand uses a very simple interface which chooses between two grips, as is found in contemporary commercial prosthetics [3]. An alternative approach is to use electromyography (EMG) electrodes to measure muscle Olaf Diegel Lund University, Lund, Sweden e-mail: olaf.diegel@design.lth.se



Fig. 1. The Touch Hand that was the first prototype designed and developed at  $\ensuremath{\mathsf{UKZN}}$ 

signals. EMG is generally done non-invasively by placing the electrodes on the skin directly above the targeted muscles. With most EMG pattern recognition systems, the user's movements are coupled directly with a grip function in the prosthetic hand [4]. EMG is the preferred method of control in commercial prosthetics as it is simple to connect the electrodes and has had the most success in prosthetic control to date. There are a variety of different EMG control techniques which each have their own strengths and weaknesses [5]. The main challenges facing EMG control are in the number of hand positions the amputee can perform. Advancements in myoelectric control will have the most direct improvement to the quality of life of amputees.

Replicating the design of the human hand is an extremely challenging engineering task. The design of a fully functional robotic hand is more complex to design as its human original. A prosthetic hand needs to have an internal power supply, control system and actuators, while still weighing the same and providing the same dexterity and grip strength as a human hand. The high cost of current prosthetic hand systems is a major limitation towards further research and adoption of new research techniques. Low-cost alternatives to prosthetics have been propelled greatly by the advancements of 3D printing, which has helped alleviate mechanical development restrictions. An example of this is the Robohand [6], greatly favoured by amputees looking for low-cost or temporary solutions, yet the Touch Hand is the only low cost hand that includes a sensory system.

This paper examines the design of the control system of the Touch Hand, seen in Fig. 1. The Touch Hand [7] is a low-cost alternative to advanced prosthetic hands, which cost anywhere from USD 30,000. The Touch Hand is being pursued to be made available for approximately USD 2,000. It is a fully dexterous hand with 6 DOFs allowing for individual control of each finger as well as rotation of the thumb. A 7<sup>th</sup> degree of freedom is located in the wrist allowing for rotation



Fig. 2. Overview of the electrical system and the interconnectivity between different units with the microcontroller unit (MCU)

of the hand. The Touch Hand is capable of all the same movements and grip types as that of the state-of-the-art prosthetic hands.

All test subjects in this paper participated voluntarily with no monetary gain, as per ethical clearance obtained. Their identities are confidential and anonymous, unless they had no objection to the photographic experimentations. The volunteers were free to refuse or withdraw from participation at any time without consequence.

#### II. ELECTRONIC DESIGN

The electrical design of the prosthetic hand control system is divided into two main areas: myoelectric (EMG) sensors and Haptic User Interfacing (HUI). These areas are shown along with the full electrical system in Fig. 2. The microcontroller unit (MCU) connects all the different aspects of the prosthetic hand and monitors the input signals from the EMG, tactile and position feedback sensors and controls the HUI, haptic feedback, and actuating circuits.

The MCU serves as the 'brain' of the prosthetic hand. The number of pulse width modulation (PWM) and analog input pins that were required was determined so that an appropriate MCU could be selected.

# A. Electromyography

EMG signals are bio-signals generated by muscles in the body when they are contracted. When a human muscle is contracted, a small potential difference is generated between the core and end of the muscle. This voltage can be read with simple surface electrodes that are placed on the skin, directly above the centre and end parts of the muscle respectively [5]. A reference electrode is also required to equate the skin's natural voltage to ground. The voltage generated this way is proportional to the extent of contraction in the muscle. This signal is very small and perfect positioning of the sensors is required to read the signal.

Most amputees retain their muscles in their residual stump, ideal for controlling the prosthesis, these muscles were used to control the hand before amputation. The basic dual-channel technique of collecting muscles signals has been kept the same as that of most contemporary upper extremity prosthetics. Two EMG sensors were used to extract the flexion and extension muscle signals from the existing forearm or from the bicep and triceps for transradial and transhumeral amputees respectively. The EMG sensors were purchased from Advancer Technologies and were selected because of their low-cost and ease of integration with Arduino circuit boards. The reference electrode is placed on a neutral/boney location such as the elbow.

In EMG prosthetics, the prosthetist conducts an EMG test with the patient to identify the best placement of the EMG sensors. The amputee's prosthetic socket is built with slots in the corresponding areas where the electrodes can be placed comfortably. This method allows for easy wearing of the prosthesis without having to first fit the EMG sensors. For testing purposes the EMG sensors were kept separate from the prosthetic socket. The signal produced by the raw EMG signal is rectified, smoothed and amplified before it is plugged into an analogue port of the microcontroller. Filtering of the EMG signal is done digitally, while the amplitude of the signal is monitored.

#### B. Haptic user interfacing

HUI refers to a method of communication between the prosthetic hand and the amputee. The HUI system uses a haptic medium (vibrotactile motors) to display information to the amputee in place of a visual display. This allows the amputee to free up his vision for other tasks and allows a more natural (haptic) communication between the prosthesis and the amputee. The HUI system uses two vibration motors to communicate with the user. These vibrators are located on the amputee's forearm with the extension vibration (referring to an upward navigation in the menu) on the top of the forearm and the flexion vibration (referring to a downward navigation in the menu) on the bottom of the forearm. This design allows the HUI vibration motors to fit inside the amputee's socket. It is possible to move the vibrotactile displays to any part of the body to allow for customizable, best-fit solutions.

#### III. CONTROL

The control system of the prosthetic hand plays two key roles. It runs and manages all the systems on the prosthetic hand (movements and positions, hand sensors, and haptic displays) and it creates a platform for the amputee to control the prosthetic hand. The ability to control a prosthetic hand has always been a major limitation in the field of prosthetics. Replacing a person's brain's natural way of controlling limbs through the nervous system is a great challenge as the method of moving one's limb is a very intuitive function. The role of the hand's control system is to simplify control while keeping it intuitive and to effectively and efficiently manage all the other systems on the prosthetic hand.

# A. Haptic user interface

The main challenge in controlling a prosthetic hand is getting it to adapt to different situations. The task of simply opening and closing the hand in a fixed grip is rather simple as the number of inputs available to the user matches the number of output states. The control of the prosthetic hand becomes significantly more complicated with each new degree of freedom introduced, as the prosthetic hand, and the number of inputs becomes less and less than the number of output states. An EMG control system that targets the muscles in the amputated arm can attempt to access up to 32 different

# 1) Navigation

The proposed system creates a HUI with the user, allowing him/her to navigate through a menu of possible grips and muscle movements and then control the selected option directly. This solution allows the user to operate an unlimited number of grips and muscle movements separately. The infinite scale of the control structure is because a list containing infinite items (of different grasps, grips, motion etc.) can theoretically be created. The size of the list does not correlate to the ability to select the first items of the list. Rather as the list size increases, the items on the end of the list



Fig. 3. (a) Overview of HUI menu; (b) Flow diagram example of HUI menu navigation



Fig. 4. Flow diagram of HUI menu navigation that the amputee could select from

muscles [8] but this is extremely costly and complicated. Commercial approaches have thus far stuck to using 2 inputs, flexion and extension. This could be flexion and extension of the wrist or forearm, depending whether the amputation is transradial (below elbow) or transhumeral (above elbow). In most amputations, targeting these muscle groups gives two inputs of variable strength that can be used to control the prosthetic hand. become increasingly more difficult to access. An overview of the menu designed for the prosthetic hand can be seen in Fig. 3.a. The HUI menu is navigated using the flex and extend inputs given by the controlling muscles. These input signals can be entered separately to navigate "up" and "down" through the menu. The flexion for the wrist has been assigned to the "down" command and extension to "up". Selection is done by inputting both signals simultaneously and the user can exit back to the "home" position by giving 2 pulses of both inputs. This "home" position can be seen as the darker blue in Fig. 3.b.

During learning the user operates the prosthetic hand in front of a computer screen which will be connected to the prosthetic hand. The screen shows the current position in the menu. This learning is done with the accompanying vibration feedback. The menu always starts in the same "home" position, allowing a sequence of navigation commands to be memorized to access a desired command. For example, the input sequence "select", "down" and "select" would activate the precision grip. Fig. 3.b shows the full menu available in the designed system. After repeated use of the system it is expected that these sequences will become very natural and act as a type of muscle memory. The gesture tab can be



Fig. 5. Logic flow diagram of pulse counting loop

expanded to contain a large variety of different gestures in a

preset format. The current menu holds seven preset grip types and four independently controllable hand gestures that give it a total of 19 unique hand positions and grip types, excluding the wrist and elbow commands.

Once a grip has been activated, the hand adjusts to the appropriate open grip position. The flex and extend signals now control the closing and opening of the hand and be sensitive to proportional values to control the grip strength. If a "back to home" command is given the navigation system will first exit to a command menu, as seen in Fig. 4. This command menu allows the user to combine useful manipulation commands like rotating the wrist or flexing the wrist or elbow. During these alternative commands the final state of the hand is maintained. For example, if the amputee wishes to pour a glass of water from a bottle, the amputee would first grip the bottle as usual. Once the bottle is gripped, the amputee can then exit to the command menu and can select the wrist rotation option. The amputee can use the wrist rotation to then pour the bottled water into a glass without dropping the bottle. The hand automatically maintains the grip on the bottle while the user controls the wrist rotation.

# 2) Feedback

The user is informed of every command by means of an appropriate vibration. The two vibration motors used for HUI relate to the two input signals that the user uses to navigate through the HUI menu. If a flexion/extension signal is read, the navigation position moves in the flexion/extension direction and the flexion/extension ("down" / "up") vibration is given indicating the successful change in menu position. The HUI feedback is only used while the user is in the menu. The sensory feedback is used during grip operation.

A summary of the navigation commands and their corresponding haptic feedback can be seen in Table I. Only four different input patterns have been taken from the two inputs in order to keep the system simple and to allow for easier tracking through the system using the haptic feedback. In future work the system could be implemented using a pattern recognition system as well, allowing quicker access to desired commands. The logic flow diagram used to extract the pulses used in the HUI navigation from the EMG signals is shown in Fig. 5.

#### 3) Using the System

The process of learning the HUI navigation menu requires

Action	<b>Flexion Signal</b>	<b>Extension Signal</b>	"Up" Vibration	"Down" Vibration
Move Up	No	Yes	Short	No
Move Down	Yes	No	No	Short
Select Option	Yes	Yes	Short	Short
Back to Home	Double Pulse	Double Pulse	Long	Long
In Grip	<b>Flexion Signal</b>	<b>Extension Signal</b>	"Up" Vibration	"Down" Vibration
Close	Yes	No	No	No
Open	No	Yes	No	No
Exit to Menu	Double Pulse	Double Pulse	Double Short	Double Short

 TABLE I:
 AVAILABLE INPUTS FOR HUI CONTROL

training. A program has been written to visually display the HUI menu on a computer screen. This program is then run in conjunction with the use of the prosthetic hand for the training period, until the user has memorised the menu. The training technique focuses on motor learning of the muscles being monitored by the EMG sensors [9]. As with all motor learning cases, motor retention is best taught through repetition of tasks



Fig. 6. Average time taken to select a grip with the HUI system

during which the tasks become more and more natural to the user. It is the goal of the HUI system approach to control that the memory of the system will migrate completely from working memory to muscle memory.

# B. Electromyography

Each amputee has differing signal levels due to the size and strength of the residual muscles, thus causing an inconsistency in the grasping motion. To account for this variation in signal strength and quality the code needs to be calibrated to each amputee's unique signals. Amplification parameters are programmed to account for differing muscle signals and filtering parameters are programmed to cancel out any interfering noise. The amplification and filtering parameters are directly linked since increasing amplification means increasing the vulnerability to noise thus more filtering is needed. Basic filtering was achieved by only recognising EMG signals that went over a preset threshold. A threshold filter helps protects against accidental triggering of commands by the amputee. Successful implementation of this control system required rigorous testing to achieve repeatable results.

# 1) Calibration and filtering

The EMG signal generated by an individual is unique. In order to correctly read an individual's EMG signal the sensors need to be calibrated and filtered. The calibration and filtering parameters need to be adjusted according to each individual. Slight adjustments are also required every time the sensors are reconnected to the muscle as the signals vary with muscle fatigue, sensor placement and arm position [5]. The EMG signal received by the micro-controller has been rectified and smoothed. This signal however still has a slight tremor as the muscles themselves do not give off a consistent voltage. The strength of the voltage signal varies greatly from person to person. To convert any extracted EMG signal into a meaningful instruction it is necessary to smooth and amplify the received EMG data.

Smoothing is done by calculating a moving average of the signal based on a RMS calculation performed on the most recent data collected in a set time frame. This time frame is referred to as the "smoothing window" of which the signal is viewed. The larger the window, the more steady and reliable the signal is, to a point, until the window starts to smooth out important data. However the larger the window, the larger the delay in response of the prosthetic hand is. The ideal window size varies from individual to individual depending on the signal received from his/her muscles. The position of the peak of the accuracy plot on the graph depends on the user's muscle tremors. A user with a larger muscle tremor requires a larger smoothing window to produce the most accurate results.

Once the correct smoothing has been performed on the EMG signal, it is necessary to map the signal so that it falls within the pre-set range expected by the controller. The user is required to perform maximum flexion tasks in order to measure the maximum signal the muscle is capable of emitting. The user is encouraged to only perform these tasks to the extent that is comfortable and easy to perform repetitively. These inputs are mapped to the standard range expected by the controller. Each muscle requires individual mapping as they will generate different amplitudes. A correctly calibrated hand allows the user to control the speed and grip force of the prosthetic hand through the changing degree of muscle flexion.

Users also differ in their ability to generate pulse values. Both the frequency and amplitude of muscle pulsing differs and needs to be individually calibrated. The user is asked to perform "pulsing" tasks with his/her muscles, and the frequency and amplitude of these pulses are measured and calibrated into the control system. Filtering is performed by measuring the amplitude of the muscle signal when the muscle is at rest. The maximum amplitude is then taken with a safety factor of 1.5 and used as the high pass filtering threshold. It was found that a typical value of 5% to 15% of the maximum flexion signal is needed to effectively filter the EMG signal. The calibration of the prosthetic hand takes around 30 minutes to perform manually for a first time user. Once the user's calibration is saved only minor calibration is required every time the sensors are reconnected.

# C. Position control

Position control of the fingers are vital to successful grasping with repeatable results. Since the fingers were actuated by means of a pulley and cable it was imperative that the controller prevent the motors from over-spooling. Over-spooling is caused when the cable slips off the pulley and entangles with the gears, causing a loss of actuation. To prevent this problem, a closed loop control was implemented over the position of the fingers. As each finger has a single degree of freedom (degree of closure), a simple flex sensor was capable of quantifying this variable for each finger. Each finger required individual calibration of its flex sensor position to set fully open and fully closed positional values. The microcontroller continually checked the value of the flex sensor and stops the finger's motor when the finger crosses the

open or closed threshold. By using the software to limit the flexion and extension of the fingers, the controller protects against unravelling of the cable spoil.

# 1) Grip control process

As the degree of closure is always known by the flex sensors, it is possible for the prosthetic hand to correctly position itself for a new grip. The position control also allows for the addition of proprioceptive sensory feedback of the fingers if desired. The grip control is done through an open loop control system on the controller closed by the user. The user is informed of the current status of the grip through the grip force and slippage haptic feedback as well as visual information from observing the hands position relative to the object. The user can then decide to close or open the hand at varying degrees accordingly.



📕 Lateral Grip 🛛 📕 Power Grip

Fig. 7. Grip force test results for lateral and power grips



Fig. 8. Hand gripping bag for passive load testing

# IV. TESTS AND RESULTS

The speed of the HUI control system was analysed. The relationship between the number of commands required to navigate the HUI menu and the time taken to access a specific grip type was determined. The physical capabilities of the hand were tested through grip force and passive loading testing. These tests were used to quantify the prosthetic hands performance. All tests conducted that required test subjects had consent and indemnity forms signed by the test subjects prior to testing, as per ethical clearance.

# A. HUI control

The Haptic User Interface is the control method developed in this research that allows the amputee to navigate through a selection of predefined prosthetic hand functions and grip types. The HUI system uses a 2-channel EMG input to give the user access to an effectively limitless number of predefined hand functions and grip types without sacrificing the ability to access or control the basic grips. In order to evaluate the effectiveness of the HUI control method a timed selection test was done to measure the time it takes to select any of the 7 predefined grips.

The test subject was an unimpaired adult male. The test was performed with a visual feedback of the HUI menu as well as with the aid of the HUI vibrotactile feedback. The subject performed 5 repetitions of each task, and the average results for each grip, with their standard deviations, are shown in Fig. 6.

# B. Physical capabilities

The prosthetic hand's physical performance is tested and discussed in this section. The ability of a prosthetic hand to grip objects is the fundamental purpose of any prosthetic hand and therefore the most critical aspect of the prosthetic hand to test. The prosthetic hand was designed to mimic the natural movements of the human hand. The shape and speed of the prosthetic hand is compared to that of a human hand to evaluate how closely this is achieved. The prosthetic hands ability to hold static loads is also evaluated as this ability determines the maximum loads that the hand can carry. The complete hand with the wrist and all motors and circuitry but without the battery weighs 540 g.

# 1) Grip force test

One of the most common measures of prosthetic hand effectiveness is done through the measurement of the grip force capable by the prosthetic hand. The power and lateral (key) grips are typically used to measure the grip force. The power grip uses all fingers and the thumb to squeeze an object and is used when gripping a bottle. The lateral grip only uses the thumb and is used to grip and handle keys, cards or other small flat objects. The hand was tested 5 times for each grip. The results of the tests are shown in Fig. 7 with their respective standard deviations.

The lateral grip of the Touch Hand had an average grip force of 3.7 N while the power grip had an average grip force of 19.5 N. The results show a large standard deviation for the power grip. The large standard deviation is due to the fact that the tips of the fingers sometimes slip on the object. The fingertip slipping on the object reduces the effective gripping force as the majority of grip force is transferred through the fingertip as this is where the tendons are connected.

# 2) Passive loading test

The passive loading test is done by gripping a bag in the hook grip position, as shown in Fig. 8. The bag is then increased in weight incrementally until failure occurs. The hand is not powered during this test and so the motors are off, thus the weight is held through the self-locking mechanism of the worm gears. The ability of the hand to hold objects without drawing power is critical in reducing the hands power consumption and extending its battery life.



Fig. 9. Amputee training with EMG control of the prosthetic hand

Weights were placed in the bag to increase the weight of the bag by 250 g increments. The hand failed at 8.25 kg with the last successful passive load being performed at 8 kg. Stretching in the tendon cables occurred before failure and caused the fingers to extend slightly. This extension did not cause the hand to drop the object and thus the cables performed satisfactorily.

# C. Practical testing

As a prosthetic hand's function is to assist an amputee with daily living and to improve his/her quality of life, the prosthetic hand needs to be tested in real life. This section investigates the prosthetic hand's ability to pick up and handle everyday objects. The hand is also tested on a transradial amputee who uses the hand to do basic tasks such as manipulating objects and drinking out of a cup.

The prosthetic hand was tested on an adult male double transradial amputee. The amputee had no experience with using EMG controlled prosthetics and owns a pair of body powered prosthetic hands. The amputee went through a period of training which lasted 30 minutes (including calibration) prior to practical use of the prosthetic hand. In Fig. 9, the amputee is viewing visual feedback of the two EMG channels on the screen of the laptop. This process helped the amputee to focus on using the correct muscles and control the degree of muscle contraction in order to control the prosthetic hand. Due to the subject's amputation, the two EMG sensors were placed over his bicep and triceps respectively. The amputee had not used these muscles properly since his amputation 20 years

prior and it took a lot of concentration to use them to control the prosthetic hand.

The amputee picked up the use of the prosthetic hand quickly and performed grasping tasks comfortably with the prosthetic hand. The amputee is shown in Fig. 10 using the prosthetic hand to drink from a cup. The amputee also was capable of pouring water from a bottle into a cup. The amputee used the prosthetic hand for 4 hours continually during this session and reported that the prosthetic hand was comfortable to use and performed satisfactorily. It was noted that the amputee's muscles experienced fatigue over the course of the 4 hour exercise. As a result of the muscle fatigue, the EMG signals weakened and the amputee felt that controlling the prosthetic hand became strenuous. Muscle fatigue would lessen with continual use of the EMG prosthetic hand.

# V. CONCLUSION

Electromyography is a tried and proven method for extracting muscle signals from a patient. The use of the residual muscles in the amputees stump is effective in mimicking the natural way in which the body controls the limbs, as the brain is used to sending signals to the muscles to control the hand. Through the suggested EMG extraction of these signals, the brain simply needs to learn the new combination of muscle movements to control the prosthetic limb.

The HUI system provides a familiar, scalable menu system that enables its users to select from a pre-set list of grip types or other hand functions while still allowing 2 DOFs of proportional control within each type. The HUI menu allows the user to directly select the desired grip without having to scroll through all the undesired grips. With the currently designed menu, the user can select from 19 unique grip and hand positions. The number of unique grip and hand position can be expanded by adding more options in the menu with marginal extra effort required from the user. The user can select from any of the 7 grip types within 5 seconds. The selection speed is satisfactory and gives the hand a good dexterity to speed ratio. The HUI system uses simple threshold control and therefore is significantly more robust than pattern recognition systems that suffer from inconsistent user muscle signals and require much longer training periods.

Training of both the controller to the user's muscle signals and of the user to correctly use the prosthetic hand is essential with the hand developed in this research. The calibration process takes about 30 minutes to complete on a first time user with only minor calibration being required after that. The position control methodology works based on a user-closed open loop methodology where feedback information is communicated directly to the user. Flex sensors in each finger are capable of accurately quantifying the fingers' positions. The flex sensors allow for software safety limits to be placed on the extension and flexion of the fingers, preventing the motors from over-unwinding the spools.

The gripping strength of the prosthetic hand is satisfactory and is capable in daily use. However, the prosthetic hands grip strength is lower than that of other contemporary prosthetic



Fig. 10. Sequence of the amputee grasping and drinking out of a cup with the prosthetic hand

hands [10]. The results from the grip stability test showed significant improvement with the use of the non-slip covering. The hand is capable of lifting suitable weights for eating and drinking purposes and meets the initial specifications. Increasing the motor power would improve grip strength and hand speed. However, a more powerful motor would be larger and heavier than the motors currently selected. Therefore a balance between size, weight and grip strength needs to be decided upon. Using more powerful motors will also increase the overall cost of the prosthetic hand.

# A. Future work

The process of identifying the ideal location for the placement of the EMG sensors can be slow and difficult. Effective software needs to be developed to aid in the process of identifying the ideal location on the amputees stump needs to be developed. This program would visually show the signal readings while the EMG sensors are placed in different locations. The calibration process currently takes up to 30 minutes to do manually. This software could also assist in the calibration process, allowing the hand to calibrate itself by running a calibration exercise with the user. This technique would also significantly reduce the calibration time for return users.

Training software would also be beneficial to the amputee to allow him/her to practice controlling the prosthetic hand with the aid of a computer program. This program will show the amputee the effectiveness of every action he/she performs. The training software will also give a visual display of the HUI interface aiding the amputee to memorize the menu layout.

# References

- [1] Carrozza, M., Vecchi, F., Sebastiani, F., Cappiello, G., Roccella, S., Zecca, M., Lazzarini, R. and Dario, P., "Experimental analysis of an innovative prosthetic hand with proprioceptive sensors," IEEE International Conference on Robotics & Automation Taipei, Taiwan, September 2003.
- [2] Wettels, N. and Loeb, G., "Haptic feature extraction from a biomimetic tactile sensor: force, contact location and curvature," IEEE International Conference on Robotics and Biomimetics, Phuket, December 2011.
- [3] Matrone, G., Cipriani, C., Carrozza, M. and Magenes, G., "Real-time myoelectric control of a multi-fingered hand prosthesis using principle components analysis," Journal of NeuroEngineering and Rehabilitation, 2012.

- [4] Jiang, N., Englehart, K. and Parker, P., "Extracting simultaneous and proportional neural control information for multiple-DOF prostheses from the surface electromyographic signal," IEEE Transactions on Biomedical Engineering, Vol. 56, No. 4, April 2009.
- [5] Castellini, C. and van der Smagt, P., "Surface EMG inadvanced hand prosthetics," Biological Cybernetics, pp. 35-47, 2009.
- [6] van As, R. and Owen, I., "Coming up short handed (the Robohand blog)," [Online]. Available: www.comingupshorthanded.com, date viewed: 13 February 2013.
- [7] van der Riet, D., Stopforth, R., Bright, G., and Diegel, O., "The Low Cost Design of a 3D Printed Multi–Fingered Myoelectric Prosthetic Hand", Mechatronics: Principles, Technologies and Applications, Nova Publishers, 2015
- [8] Khushaba, R., Kodagoda, S., Takruri, M. and Dissanayake, G., "Toward improved control of prosthetic fingers using surface electromyogram (EMG) signals," Expert Systems with Applications, Vol. 39, pp. 10731-10738, 2012.
- [9] van der Riet, D., Stopforth, R., Bright, G. and Diegel, O., "Simultaneous vibrotactile feedback for multi-sensory upper limb prosthetics," 6th IEEE Robotics and Mechatronics Conference of South Africa, Durban, October 2013.
- [10] van der Riet, D., Stopforth, R., Bright, G. and Diegel, O., "An overview and comparison of upper limb prosthetics," IEEE AFRICON Conference, Mauritius, September 2013.

# Contrasting Classifiers for Software-based OMR Responses

Bertram Haskins School of Information and Communication Technology Nelson Mandela Metropolitan University Port Elizabeth, South Africa, 6001 Email: Bertram.Haskins@nmmu.ac.za

*Abstract*—Systems based on optical mark recognition (OMR) provide a means to address rising student numbers in university modules. Developing a system which addresses the needs of a specific module provides added flexibility to enhance the teaching environment. The process is complicated by responses which have been selected and then deselected by the respondent.

This study contrasts four classifiers for identifying selected responses on OMR answer sheets. Four classifiers are constructed based on features derived from the number of pixels in an image, patterns derived from the image, edit distances derived from bit string patterns and the average number of edges per axis. The classifiers based on the number of pixels and edges make use of simple thresholds for classification, whereas the classifiers based on patterns and edit distance make use of classification trees.

The classifier based on the number of pixels in the image delivers the best results as it has very high levels of accuracy, sensitivity and specificity, but is unable to identify any responses which have been deselected by a respondent. The classifier based on the number of edges per axis performs very well in accurately classifying positive responses, but has only moderate levels of general accuracy. The edge-base classifier is, however, the only classifier able to correctly classify any deselected responses.

Index Terms—Image Processing, Image Edge Detection, Decision Trees

# I. INTRODUCTION

Rising student numbers have been both a blessing and a burden in recent years. The increase in student number caps per subject have led to more students being allowed to university, but have necessitated a re-evaluation by many lecturers as to the management and administration of their subjects.

One solution is the introduction of more multiple choice assessments into either the course work of the subject or its final examination. In a day-to-day classroom situation offline (i.e. paper-based) multiple choice assessments provide a lecturer with more freedom for conducting class tests in theory classes, without the requirement of a PC for every student. Although there are many tools available for creating off-line multiple choice tests or answer sheets, such as Auto Multiple Choice (AMC) [1], these tools may either require a different operating system or may not be as flexible as required.

# II. PROBLEM STATEMENT AND OBJECTIVES

In order to create an off-line multiple choice system, based on optical mark recognition (OMR), which allows a user to vary the layout, number of hotspots (response areas) and size of the hotspots; a flexible software-based system was created. This system consists of a computer-based interface for generating student answer sheets and marking scanned, completed answer sheets. In order to mark an answer sheet, the system has to be able to identify which of the hotspots have been selected. This task may not be as simple as it seems, as a student may not follow the instructions for completing the answer sheet or may select a specific hotspot and the deselect it by drawing a line through it. Although it may be simpler to penalise a student for non-adherence to the rules of the answer sheet, there is a big push towards humanising pedagogy at the university at which this study is conducted. Humans tend to make mistakes; especially in a pressure-filled situation such as a test or exam. Therefore, the main objective of this study is to determine which classification technique is most adept at identifying both selected and deselected hotspots.

The remainder of this study is structured as follows; Section III provides an overview of the system on which this study is based. Section IV delves into the various approaches used to facilitate the OMR process in this study. Section V discusses the results of implementing the various OMR approaches. Finally, Section VI provides some insight into the utility of the approaches followed in this study and proposes a future direction for the research.

# III. BACKGROUND

As the name suggests, optical mark recognition is the process of detecting marked responses on an answer sheet. This is generally done by a device known as an optical mark reader, which identifies marks by identifying those areas which register significantly less light than the surrounding paper [2]. OMR need not always involve a specialised optical mark reader. Some approaches, such as AMC [1], queXF [3] and SQS (Shared Questionnaire System) [4], may be classified as OMR software, in that they make use of a normal desktop scanner and computer software to determine which areas on a predefined answer sheet have been marked.

Although the university at which this study is conducted provides an optical mark reader-based approach, it has a fixed number of responses per item. It therefore does not lend itself well to multiple choice multiple answer or even match column A to column B type questions; as they would both require more than the standard 4 or 5 hotspots per question. To meet these requirement it was decided to create a more flexible softwarebased OMR application to aid in both the creation, marking and management of multiple choice-style assessments.

The application allows for a multiple choice assessment to be created using a simple script; without the need to know a typesetting language such as LaTeX (used by AMC). The script allows the user to set different question areas on the answer sheet and control where the question areas should appear on the sheet. For individual sub-questions, the examiner simply needs to specify the question label, total for the subquestion, value of a single correct entry and a series of letters to print on the answer sheet as possible answer boxes. An answer is marked as the correct one by including the number *I* directly adjacent to its representative number. Furthermore, the totals of the individual questions, as well as the total for the paper is automatically calculated from the script.

The script is used to automatically generate an answer sheet (an example of which is shown in Figure 1) and memorandum. An answer sheet must be printed for every student and then scanned and provided to the application as input after the test. Each answer sheet contains an area for a student to write their student number, name and surname. Hotspots are provided for the student to enter his or her student number. The system allows for the creation of *practice* answer sheets on which a student can write their initial answers before completing the final version. When conducting a test, a student is provided with an extra practice answer sheet (labelled with a watermark) and explicitly told that only properly completed hotspots on the main answer sheet will be marked. Crossing out a hotspot to deselect it once selected is not permitted. Even with the extra answer sheet and instructions, there are still students who do not adhere to the instructions not to cross out a hotspot once completed. Furthermore, students tend to draw lines and arrows or write alternative answers (letters) on the answer sheets. For this study, these extra lines, arrows and text will be ignored as the approach to process them does not fall within the scope of this study.

# IV. IMPLEMENTATION

The system devised for this study is designed to process a series of scans of completed multiple choice answer sheets. Every sheet is processed independently and compared to a memorandum. The scanning of the tests is an external process and does not form part of the marking system. For purposes of clarity, the process required for marking a single student answer sheet is graphically illustrated in Figure 2.

Once an individual image of an answer sheet has been loaded it first needs to be de-skewed. The system makes use of the rectangle around the answer sheet to calculate the skew-gradient and then counter-rotate the image to counter the skew. Once the image has been de-skewed, the system makes use of this same rectangle to resize the image to conform to the size of the memorandum. Finally, the image is converted to a greyscale copy. After these three initial steps, the memorandum is consulted to determine where the



Figure 1. A Sample Answer Sheet



Figure 2. Process Flow in the OMR-Based Marking System

relevant hotspots are situated on the image. Each hotspot is then processed in turn.

The system was developed with the ability to vary the size of the hotspots. For this study, the hotspots were set to a fixed size of  $50 \times 50$  pixels. Because of the slight variances which might still exist in the scanned images, regardless of efforts to de-skew and resize the image, every hotspot is read into memory with a total padding of 11 pixels to each side of the hotspot. This yields a final hotspot image bounding box of  $72 \times 72$  pixels. This extra padding, captured as part of the hotspot, is not only used as a counter to the skew and size of the image, but also to capture information regarding the area around the hotspot.

In many cases students may change their mind regarding their selection of a specific hotspot. When doing so, they tend to try and cross out the hotspot and then select another. Figure 3 shows a few examples of crossed out hotspots. In all of these cases there is either a secondary layer of marking over the selected hotspot or marks around the hotspot,



Figure 3. Examples of Crossed Out (Deselected) Hotspots

Table I					
Data	Set	Detail			

Description	Training	Test
Number of files	74	60
Total number of hotspots	24255	11640
Number of Selected hotspots	3725	2930

which indicate that the hotspot was initially selected and then deselected.

#### A. Source Data

The system was developed and tested using two multiple choice-based tests from two different subjects. The scanned answer sheets were saved as JPEG files, with file names based on sequential numbering. No student identification details were used as part of this study. For both of these groups, the answer sheets were manually validated and compared with the memorandum. The hotspots were labelled and referenced according to file name and whether or not a hotspot was selected (Y) or not (N). The files from one group serve as the training set for this study and the files from the other group serve as a test set. The details of both data sets are contrasted in Table I. All of the entries in these two data sets are clearly selected or not selected. Care was taken to avoid entries which have been selected and then deselected.

A third data set, consisting of only 20 entries was created to serve as a validation data set. This very small data set consists of entries which may confuse the chosen classifier. In many cases respondents mark a hotspot, change their mind and fill in another hotspot and then simply cross out the original hotspot. Examples of these hotspots are shown in Figure 3. These entries are simple for a human marker to spot, but might be more difficult for an automated classifier. All the entries in this data set have been identified, so that a classifier should want to classify them as selected. However, in each entry there is extra information layered across or surrounding the hotspot which serves to indicate that this selection has been deselected. Thus, all entries in this data set have been labelled as not selected (N). These 20 entries come from the same tests used as input data for the training and test data sets and were the only deselected entries across all of the selected hotspots, thus constituting only around 0.003% of all selected hotspots. Although this number may seem inconsequential, the low number may be the result of explicit instruction to the respondents, which might not occur if an exam is overseen by an independent overseer, and the availability of a second answer sheet which serves as a practice sheet.

# B. Feature Selection

In order to identify which hotspots have been selected or not, it is necessary to extract features from the  $72 \times 72$ 



Figure 4. Hotspot Size Reduction

pixel image. Features may be described as the input values to a selection process. Selecting the correct features is one of the main steps in the data analysis process and is a key contributing factor in the success of any machine learning endeavour [5, p. 4].

To this end four feature separate feature sets were devised to describe each hotspot. The first feature set is based on a brute force pixel count. The number of pixels part of the selection is determined by calculating the average RGB value for each pixel. If the average colour value of the pixel is higher than 200 it is flagged as selected. The average value of 200 was arrived at by experimentation. For every hotspot a percentage is calculated of how many of the total number of pixels in the hotspot have been selected. This feature set is referred to as PIXELS.

For both the second and third feature sets, the hotspots are processed by an algorithm which serves to decrease the number of possible features by shrinking the size of the hotspot. A hotspot is iteratively processed until its size is reduced to  $9 \times 9$  pixels. On each iteration the size of the hotspot is halved. The target hotspot is processed in blocks of 4 pixels ( $2 \times 2$ ) each; with each 4 blocks on the input image being represented as a single block on the output image. This has the effect of reducing the input image to a quarter of its input size on each iteration, as shown in Figure 4. If at least 3 of the 4 processed pixels are selected, then this area is indicated as selected on the new, smaller image.

This smaller,  $9 \times 9$  pixel image is used to generate the second and third feature sets. For the second feature set, entitled PATTERNS, the reduced hotspot image is processed in blocks of  $3 \times 3$  pixels, as illustrated in Figure 5. Initially every hotspot in the training data was subjugated to this process to determine which patterns of 9 pixels occur within the hotspots. Each of these patterns can be written as a bit string which is 9 digits in length. This conversion is also illustrated in Figure 5. Using this process, a  $9 \times 9$  pixel hotspot can then be converted into 9 separate bit strings.

Using the pattern library as a filter, every bit string is compared to the library to determine if the pattern is one associated with a selected hotspot. If the pattern corresponds to a known bit string pattern, its value is kept; otherwise it is replaced with a bit string which consists of all zeroes. Once this process is complete, each bit string is converted to its decimal equivalent. This results in 9 decimal val-



Figure 6. EDGE Feature Extraction Process

ues, each in the range 0 to 512, to act as input features for the PATTERN\_VALUES set. The 9 features in this set are labelled TOP\_LEFT, TOP\_CENTER, TOP\_RIGHT, MID-DLE\_LEFT, MIDDLE\_CENTER, MIDDLE\_RIGHT, BOT-TOM\_LEFT, BOTTOM\_CENTER and BOTTOM\_RIGHT.

The third feature set, BIT\_STRINGS, describes a hotspot as a single bit string that is 81 bits in length. This bit string is created by concatenating the PATTERNS features in the numerical sequence specified in Figure 5.

The fourth feature set, EDGES, consists of a single number describing the average number of edges encountered per axis if all of the pixels on a hotpspot are traversed horizontally and vertically through the middle of the hotspot. Various built-in methods, included in the AForge.NET Framework [6], were used to reduce each hotspot to a series of simplified edges. The process involves several steps.

Initially a hotspot is smoothed, using bilateral smoothing, in order to remove unnecessary noise in the image, but still preserve any edge information. In the second step, the image is converted to binary (i.e black and white, 8 bits per pixel), from the input greyscale. In the third step a Canny edge detector is applied to the image to highlight any edges. The image is then colour inverted to serve as input for the next step which removes any black holes in white portions of the image which are smaller than  $15 \times 15$  pixels. The final step in the process counts the number of edges on the image along the central horizontal and vertical axes to determine the average number of edges per axis. This entire process is graphically represented, along with examples of selected, not selected and deselected hotspots, in Figure 6.

#### C. Detection Approaches

The OMR software created for this study is structured in a modular fashion, so that key processes may be easily replaced with equivalent ones. This section details four alternative



Figure 7. Classification Tree Derived From the PATTERN VALUES Features of the Training Data

approaches to the *Classification* process as highlighted in Figure 2.

1) Counting Pixels: This classifier makes use of the PIX-ELS feature set. By applying a simple threshold to the percentage value in the feature set, a hotspot may be flagged as selected or not. The threshold for selection is set at 20%. This threshold was found by determining the percentage of selected pixels of each of the hotspots in the training set which were manually labelled as not selected. From this process it was determined that the maximum number of selected pixels for an unselected image is around 20%. The strength of this approach is that it is not affected by any variance in the location of the hotspot within the bounding box which may be the result of slight skew or size difference. This process may be seen as the equivalent of an optical mark reader contrasting the light and dark areas on a page.

2) Comparing Patterns: This process makes use of the PATTERNS feature set. Using the training data set as input, a decision tree was trained on the 9 features contained in the feature set. A decision tree is a branching structure which serves as a representation of a series of logical decisions. Each decision point is modelled as a node, which corresponds to an input variable from a data set. The branches extending from each node represent the results of each decision.

R is a language and environment for statistical computation and graphical representation [7]. It is a GNU project which provides a wide variety of statistical and graphical techniques. The *rpart* package in R allows for the creation of classification or regression trees [8]. To create a classification tree, the output feature must be a nominal value, whereas regression trees may be built using ordinal data. For this study the *rpart* package was used to create a classification tree; using the default settings of the package.

The classification tree derived from the training data is shown in Figure 7. Every node split illustrates the decision made at the specific node. The leaf nodes specify which specific response value constitutes the majority of the population and the sub-text of each leaf node (n = x) specifies the total population of the training set which meets the leaf node criteria.



Figure 8. Scaled Tree Variable Importance for PATTERNS

Figure 8 illustrates a scaled variable importance diagram for the features in the training data set. In essence, a variable importance diagram illustrates the overall importance of a specific variable in support of the splitting decisions made at every node, either as a primary or a surrogate splitter. The primary splitter is the variable on which the decision has been made, but a surrogate splitter is a variable which may be used to achieve similar results as when using the primary splitter. From this diagram it can be seen that the main contributing features are those of MIDDLE\_CENTER, MIDDLE\_RIGHT, BOTTOM\_CENTER AND BOTTOM\_RIGHT. This is also supported by the decision tree in Figure 7 as these are the features on which all of the branching decisions are made.

3) A Text-Based Approach: This process makes use of the BIT\_STRINGS feature set. Each entry in this sets consists of a single 81-bit bit string. In order for the text-based approach to work it was necessary to first determine a base bit string which serves as a representation of a selected hotspot. To this end, the training data set was processed to select only the entries in the set which have been pre-labelled as selected. Each of these 81-bit bit strings was compared on a bit-by-bit basis to determine which bit value (0 or 1) was most prevalent for each position in the string. This process yielded an average bit string value to be used in a string comparison process.

Levenshtein proposed a method to measure the cost involved in transforming one string into another [9]. This is also sometimes called the edit distance and is a measurement of the number of insertions, deletions and substitutions necessary for the transformation to occur. Levenshtein distance is used in many modern applications for such diverse purposes as keyword extraction from text data [10] and normalising microtext [11].

Each of the entries in the training BIT STRINGS feature set was compared to the predetermined *average* bit string to determine the edit distance between these two strings. The comparison took place in segments of 9-bits each, with each 9 bit string corresponding to a region on the original input hotspot. This process yielded 9 separate edit distance measurements, per input hotspot, which were used to train a classification tree in the same fashion as described in Section IV-C2. This approach was selected as it was thought that it would highlight areas in which a hotspot differs from



Figure 9. Scaled Variable Importance for BIT STRINGS

a generalised form. These differences may indicate non-selection or deselection.

Figure 9 illustrates the scaled variable importance diagram for the features in the training data set. For this feature set the most important contributing features are MID-DLE\_CENTER, BOTTOM\_CENTER, MIDDLE\_RIGHT and BOTTOM\_RIGHT. This shows a difference with regards to the variable importance of the approach based on the PATTERNS. The PATTERNS features seem to indicate that the entire bottom-right quadrant of the image is important whereas the BIT STRINGS seem to suggest that the middle of the image is the most important to consider when making a classification decision.

4) Counting Edges: This process makes use of the EDGES feature set. Through close scrutiny of the EDGES features, it was experimentally discovered that hotspots which are not selected have an axis edge count average of less than 2. The processing steps described in Figure 6 tend to yield an empty image or one with very few edges for hotspots which have not been selected. Thus, a simple classifier was constructed which classifies hotspots with an axis edge count average of 2 or higher as selected and any others as not selected. This approach was chosen as the average number of axis edges would not be as heavily influenced by the precise location of the hotspot on the image, as would the features identified for the PATTERNS and BIT\_STRINGS feature sets.

#### V. RESULTS AND FINDINGS

The features generated from the training data set were used to develop the classifiers. In order to test the performance of the classifiers it was necessary to test them on data unrelated to the training process. Thus, all four classifiers were applied to features generated from the test data set.

The results of the data sets were contrasted using 4 metrics, namely accuracy, sensitivity, specificity and Kappa. Accuracy is a measure of the number of true results among all of the cases examined, sensitivity is a measure of the rate of true positives and specificity is a measure of the rate of true negatives. All three these values are expressed as values in the range 0 to 1. Values approaching 1 are considered better results.

Cohen's Kappa is a statistical calculation for summarising the cross-classification of two nominal variables [12]. Total

Table II Comparison of Accuracy, Sensitivity, Specificity and Agreement (kappa) of the Four Classifiers

Metric	PIXELS	PATTERNS	BIT STRINGS	EDGES
Accuracy	0.9998	0.9605	0.9380	0.7357
Sensitivity	1.000	0.8911	0.7760	0.9450
Specificity	0.9998	0.9838	0.9924	0.6653
Kappa	0.9995	0.8929	0.8235	0.4651

agreement between classifiers, is signified as 1. Zero signifies the expected result under total independence and a negative value occurs when the calculated agreement is less than what would be expected under chance conditions. In this study Kappa is calculated by comparing the results generated by a classifier with the pre-classified label of each data point.

The results of these calculations, as shown in Table II, indicate that the classification task might not be a difficult one, as all the classifiers yield very good or acceptable results. The best results are obtained from the simplest of all the classifiers, namely a simple pixel count and a threshold. This classifier was able to correctly classify all the selected hotspots and misclassified only two unselected hot spots. The two classifiers based on decision trees both have a very high level of accuracy, but the classifier based on the BIT\_STRINGS feature set has a lower sensitivity at 0.7760. This suggests that it performs slightly worse at identifying positive cases (selected hotspots). On accuracy, specificity and Kappa, the classifier based on the EDGES features performed the worst of all four classifiers, but was only outperformed by the PIXELS classifier with regards to sensitivity. This means that it is very adept at classifying selected hotspots.

There is however an issue with regards to multiple choice responses which may confuse the classifiers; namely, determining if a selected hotspot was deselected. Therefore, the validation data set was used to determine whether any of the four classifiers would be successful at correctly labelling these confusing entries. From the 20 entries in this data set the EDGES classifier was able to correctly classify 45% of the entries. None of the other classifiers were able to correctly classify any of the entries in the data set.

#### VI. CONCLUSION AND FUTURE WORK

The main objective of this study was to determine which classification technique is most adept at identifying both selected and deselected hotspots. To this end, four classifiers were created based on four different feature sets. The simplest of the classifiers performed the best as it had the highest levels of accuracy, sensitivity, specificity and agreement (Kappa).

However, this classifier yielded very poor results at identifying hotspots which have been selected and then deselected. As this classifier is based on a simple count of the number of pixels in the image, it does not take into account what the different pixels represent. Another classifier, based on the number of edges per axis, seemed to perform poorly on general classification, in contrast to the other classifiers, but outperformed all the other classifiers with regards to identifying hotspots which have been selected and then deselected. It was able to correctly classify 45% of the hotspots on which it was tested, whereas none of the other classifiers were able to correctly classify any of the entries.

Thus, to satisfy the main objective of this study it is recommended that a two-fold classification process be employed. The classifier based on the pixels would be used to identify selected hotspots. After this process, the classifier based on the average number of edges per axis would be applied only to those hotspots classified as selected, to determine if the hotspot was deselected.

The results of this study have shown that, in principle, it is possible to automatically determine whether hotspots have been deselected. This might serve as a possible life-line to students when making mistakes. Such a process might also serve as an aid to lecturers for improving the accuracy of automated marking and averting unnecessary student queries. The study has also shown that more complicated processes, such as ones based on edge detection, do not necessarily perform better at rudimentary tasks.

Although the number of hotspots which have been deselected are very few in comparison to the rest; this may be the result of proper instruction to the respondents and the availability of a second answer sheet. A future study will focus solely on these types of entries and any other information on the answer sheets, such as notes, lines and arrows created by the respondent, to try and improve the classification rate of these entries.

### REFERENCES

- AMC, "Multiple choice sheets automated marking," Web, 2015, retrieved: 17 March 2015, http://home.gna.org/auto-qcm/.
- [2] ITL Education Solutions Limited, Introduction to Information Technology. Pearson Education, 2005.
- [3] Australian Consortium for Social and Political Research Incorporated, "queXF Overview," Web, 2015, retrieved: 3 September 2015, http://quexf.sourceforge.net/node/1.
- [4] Jean-Philippe Lang, "Shared Questionnaire System Overview," Web, 2015, retrieved: 3 September 2015, http://dev.sqs2.net/projects/sqs.
- [5] I. Guyon, S. Gunn, M. Nikravesh, and L. Zadeh, *Feature Extraction: Foundations and Applications*, ser. Studies in Fuzziness and Soft Computing. Springer Berlin Heidelberg, 2006.
- [6] AForge.NET, "AForge.NET Framework," Web, 2015, retrieved: 2 September 2015, http://www.aforgenet.com/framework/.
- [7] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria, 2015. [Online]. Available: http://www.R-project.org/
- [8] T. Therneau, B. Atkinson, and B. Ripley, *rpart: Recursive Partitioning and Regression Trees*, 2015, r package version 4.1-9. [Online]. Available: http://CRAN.R-project.org/package=rpart
- [9] V. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," *Soviet Physics Doklady*, vol. 10, no. 8, pp. 707–710, 1966.
- [10] T. Runkler and J. Bezdek, "Automatic keyword extraction with relational clustering and Levenshtein distances," in *Proceedings of the 9th IEEE International Conference on Fuzzy Systems*, vol. 2, IEEE. IEEE, 2000, pp. 636–640.
- [11] B. Haskins and R. Botha, "A mixed-method apporach to normalising Dr Math microtext," in *Proceedings of the International Conference on Computational Science and Technology (ICCST)*. Institute of Electrical and Electronics Engineers, 27 – 29 August 2014.
- [12] M. J. Warrens, "Cohen's kappa is a weighted average," *Statistical Methodology*, vol. 8, no. 6, pp. 473 484, 2011.

# Switching Cascade Controllers Combined with a Feedforward Regulation for an Aggregate Actuator in Automotive Applications

Paolo Mercorelli, Nils Werner, Theo van Niekerk, Leander Behre and Udo Becker

*Abstract*— This work is devoted to a hybrid actuator which consists of a piezo and a hydraulic structure. Moreover, a cascade controller which consists of an internal PID and an external one combined with a feedforward regulator is proposed for the control strategy in an engine application. The main idea of this contribution is to realize a camless engine with its control. Simulation results consist of tracking of exhaust valve trajectories are shown.

Key-words: PID controllers, Lyapunov's approach, piezo actuators

# I. INTRODUCTION

Figure 1 indicates a possible new engine structure with, evidently, four piezo actuators. Generally speaking, this project closely concerns a new conception of the functionality of some parts of the engine. In particular, a camless engine is proposed. The idea which this paper presents is to use hybrid actuator



Fig. 1. New structure of the engine.

composed by a piezo part and a hydraulic one in order to take advantages of both: the high precision and velocity of the piezo part and the force of the hydraulic one. Hybrid actuators represent a viable solution to find some compromise for control systems specifications such as precision, velocity and robustness.

In this paper the hysteresis effect is a model using a linearization. A linear boundary of the hysteresis is considered

and a switching approach is used to follow the hysteresis characteristics. The easiest idea is to consider the upper and the lower bound of the linear characteristic. PID regulators are very often used in industrial applications because of their simple structure, even though in the last years advanced PID controllers have been developed to control nonlinear systems.

The objective of this paper is to show:

- A model of the a hybrid actuator
- A PID-PID cascade regulator combined with a feedforward regulator.

The advantage of using PID controller is due to their easy practical implementation. Control PID structure is a switching one which considers the proposed switching model of the hysteresis effect in order to eliminate the hysteresis effect on the control output. Hysteresis effects are well known and different proposed compensation schemes are present in practical applications such as that proposed in [1].

The paper is organized with the following sections. Section II is devoted to the model description. In Section III an algorithm is shown to derive the PID control laws. The paper ends with Section IV in which simulation results of the proposed valve using real data are presented. After that the conclusions follow.

# II. MODELING OF THE PIEZO HYDRAULIC ACTUATOR

In the diagram of Fig. 3 the T-A connection links the couple of valves with the tank and the P-B connection links the couple of valves with the pump.

Figure 2 shows in detail a part of the hybrid structure which consists of a piezo actuator combined with a mechanical part. These two parts are connected by a stroke ratio to adapt the stroke length. The proposed nonlinearity model for PEA is quite similar to these presented in [2] and in [3]. This model is constructed from a sandwich model as shown in Fig. 4. Figure 5 shows the equivalent circuitry for a PEA with the I-layer nonlinearities of hysteresis and creep, in which two I-layers are combined together as  $C_a$  and  $R_a$ . The capacitor,  $C_a$ , is an ordinary one, which might be varied slightly with some factors, but here it would be assumed constant first for simplicity. The resistor,  $R_a$ , however, is really an extraordinary one with a significant nonlinearity. The resistance is either fairly large, say  $R_a > 10^6 \Omega$ , when the voltage  $||V_a|| < V_h$ , or is fairly small, say  $R_a < 1000$ , when  $||V_a|| > V_h$ . In [3], the threshold voltage,  $V_h$ , is defined as the hysteresis voltage of a PEA.

Based on this proposed sandwich model and the equivalent circuitry as shown in Fig. 5, we can further derive the state

Paolo Mercorelli is with the Institute of Product and Process Innovation, Leuphana University of Lueneburg, Volgershall 1, D-21339 Lueneburg, Germany. Tel.: +49-(0)4131-677-5571, Fax: +49-(0)4131-677-5300. mercorelli@uni.leuphana.de. Nils Werner and Udo Becker are with the Faculty of Automotive Engineering, Ostfalia University of Applied Sciences, Kleiststr. 14-16, D-38440 Wolfsburg, Germany. Tel. +49-(0)5361-831615 Fax. +49-(0)5361-831602. n.werner@ostfalia.de Theo van Niekerk and Leander Behre are with the Department of Mechatronics, Nelson Mandela Metropolitan University, Summerstrand, 6031 Port Eliszabeth, South Africa.


Fig. 2. Structure of the hybrid actuator and detail of the slits



Fig. 3. Scheme of the whole Hybrid Piezo Hydraulic structure



Fig. 4. The sandwich model of the PEA



Fig. 5. Electrical part of the model

model as follows:

$$\dot{V}_{a}(t) = -\left(\frac{1}{R_{a}} + \frac{1}{R_{o}}\right)\frac{V_{a}(t)}{C_{a}} - \frac{V_{z}(t)}{C_{a}R_{o}} + \frac{V_{in}(t)}{C_{a}R_{o}} \quad (1)$$
  
$$\dot{V}_{z}(t) = \frac{\dot{Q}_{b}}{C_{z}} + \frac{1}{C_{z}}\left(-\frac{V_{a}(t)}{R_{o}} - \frac{V_{z}(t)}{R_{o}} + \frac{V_{in}(t)}{R_{o}}\right), \quad (2)$$

where  $Q_b = D_y F_z(t)$  is the "back electric charge force" (back-ecf) in a PEA, see [3]. According to [3] and the notation of Fig. 7, it is possible to write:

$$F_z(t) = M_p / 3\ddot{x}(t) + D\dot{x}(t) + Kx(t) + K_x x(t).$$
(3)

K and D are the elasticity and the friction constant of the spring which is antagonist to the piezo effect and it is incorporated in the PEA.  $C_z$  is the total capacitance of the PEA and  $R_o$  is the contact resistance. The following notation is used:  $K_x$  is the elasticity constant factor of the PEA. Factor  $D_x K_x = T_{em}$  has the name "transformer ratio" in the scientific literature and the most important characteristic of the electromechanical transducer is described by it.  $M_p/3$  being in our case the moving mass of the piezo structure is a fraction of whole piezo mass.  $M_{SK}$  is the sum of the mass of the piston with the oil and the moving actuator and  $M_v$  is the mass of the valve. The moving mass of the piezo structure is just a fraction of the whole piezo mass. The constructor of the piezo gives the value of this fraction and it is determined by experimental measurements.  $K_{SK}$  and  $D_{SK}$  are the characteristics of the antagonist spring to the mechanical servo valve, see Fig. 7.  $D_{oil}$  is the friction constant of the oil. Moreover, according to [3], motion  $x_p(t)$  of diagram in Fig. 6 is

$$x_p(t) = D_x V_z(t). \tag{4}$$

According to the diagram of Fig. 5, it is possible to write as follows:

$$V_z = V_{in}(t) - R_0 i(t) - H(x_p(t), V_{in}(t)),$$
(5)

where  $R_0$  is the connection resistance and i(t) is the input current as shown in Fig. 5.  $H(x_p(t), V_{in}(t))$  is the function which describes the hysteresis effect mentioned above and shown in the simulation of Fig. 6. Considering the whole



Fig. 6. Simulated Hysteresis curve of the piezo part of the actuator:  $H(x_p(t), V_{in}(t))$ 

system described in Fig. 7, the electrical and mechanical systems described in Figs. 5, 6 and 7 can be represented by the following mathematical expressions:

$$\frac{M_p}{3}\ddot{x}(t) + M_{SK}\ddot{x}_{SK}(t) + Kx(t) + D\dot{x}(t) + K_{SK}x_{SK}(t) + D_{SK}\dot{x}_{SK}(t) + D_{oil}\dot{x}_{SK}(t) + K_x\Big(x(t) - \Delta x_p(V_{in}(t))\Big) = 0, \quad (6)$$

where  $\Delta x_p(t)$  represents the interval function of  $x_p(t)$  as shown in Fig. 6 which, according to equation (4), can be expressed as:

$$\Delta x_p(t) = D_x \Delta V_z(t). \tag{7}$$

Finally, using equations (5) and (7),

$$K_x \Delta x_p(t) = K_x D_x \big( V_{in}(t) - R_0 i(t) - H(\Delta x_p(t), V_{in}(t)) \big),$$
(8)

which represents the interval force generated by the piezo device. Equation (6) can be expressed in the following way:

$$\frac{M_p}{3}\ddot{x}(t) + M_{SK}\ddot{x}_{SK}(t) + Kx(t) + D\dot{x}(t) + K_{SK}x_{SK}(t) + D_{SK}\dot{x}_{SK}(t) + D_{oil}\dot{x}_{SK}(t) + K_xx(t) = K_x\Delta x_p(V_{in}(t)).$$
(9)

It is to be noticed that the following relationship holds:

$$x_{SK}(t) = Wx(t), \tag{10}$$

where W is the position ratio above defined and it states the incompressibility of the oil in the conic chamber. The following equation completes the dynamic of the considered system:

$$M_v \ddot{y}(t) + D_{oil} \dot{y}(t) = F(x_{SK}(t), y(t)) - F_d(t).$$
(11)

According to Fig. 7, x(t) is the position of the piezo actuator,  $x_{SK}(t)$  is the position of the mechanical servo actuator, y(t) represents the position of the valve. Function  $F(x_{SK}(t), y(t))$  represents the force exerted by the pump on surface S of the armature of the moving valve, see Fig 7. Moreover,

$$F(x_{SK}(t), y(t)) = (p_A(t) - p_B(t))S,$$
(12)

where  $p_A(t)$  and  $p_B(t)$  are the pressure in the two oil chambers separated by the armature of the valve. In [4] pressure  $p_A(t)$ and  $p_B(t)$  are nonlinear functions of the mechanical servo valve position y(t). The linearization of this function at each desired position  $y_d(t)$  will be considered later.  $F_d(t)$  is the combustion back pressure in terms of force. According to Fig.



Fig. 7. Model of the whole actuator

6 in which an upper bound and a lower bound of the hysteresis curve are indicated, it is possible to write that:

$$\Delta x_p(V_{in}(t)) = \begin{bmatrix} -a & a \end{bmatrix} + bV_{in}(t), \tag{13}$$

with  $a \in \mathbb{R}$  and  $b \in \mathbb{R}$  two positive constants are indicated. In particular,

$$\underline{\Delta}x_p(V_{in}(t)) = -a + bV_{in}(t), \qquad (14)$$

and

$$\overline{\Delta}x_p(V_{in}(t)) = a + bV_{in}(t). \tag{15}$$

Considering this notation, the system represented in (6) can be split into the following two systems:

a)  

$$\frac{M_p}{3}\ddot{x}(t) + M_{SK}\ddot{x}_{SK}(t) + Kx(t) + D\dot{x}(t) + K_{SK}x_{SK}(t) + D_{SK}\dot{x}_{SK}(t) + D_{oil}\dot{x}_{SK}(t) + K_xx(t) = \underline{\Delta}x_p(V_{in}(t)),$$
(16)

b)  

$$\frac{M_p}{3}\ddot{x}(t) + M_{SK}\ddot{x}_{SK}(t) + Kx(t) + D\dot{x}(t) + K_{SK}x_{SK}(t) + D_{SK}\dot{x}_{SK}(t) + D_{oil}\dot{x}_{SK}(t) + K_xx(t) = \overline{\Delta}x_p(V_{in}(t)),$$
(17)

#### III. A SWITCHING CASCADE PID-PID CONTROLLER COMBINED WITH A FEEDFORWARD REGULATOR

Figure 8 shows the proposed control structure. It is possible to see how a feedforward structure is present in order to achieve a regulation around the desired trajectory. An internal PID structure and an external one complete the whole control scheme to guarantee the robustness. The feedforward control structure is based on the inversion of equations (16) and (17) together with equation (10) in which just the mechanical part of the model is considered because the piezo part is faster that the mechanical one. The following expression state this inversion:

$$u_{feed}(t) = \frac{\left(\frac{M_p}{3} + M_{SK}W\right)\ddot{y}_d(t)}{K_x b} + \frac{\left(D + (D_{SK} + D_{oil})W\right)\dot{y}_d(t)}{K_x b} + \frac{(Ky + K_x + K_{SK})y_d(t) + K_x(-1)^q a}{K_x b}.$$
 (18)



Fig. 8. Proposed control structure

#### A. A switching internal PID controller

As already mentioned, the PID control structure presented in this section is quite similar to the sliding control structure presented in [5]. In [6] an integral action is considered in a similar way as in this contribution in order to improve the controller performance in steady state (zero error). The proposed integral sliding mode control is derived using Lyapunov approach. Lyapunov approach is very used to derive control law also in industrial applications.Considering Figs. 5 and 7, if the dynamics shown in (16) and (17) are considered in a state space representation, then

$$\dot{x}_1(t) = x_2(t)$$
 (19)

$$\dot{x}_{2}(t) = \frac{-Dx_{2}(t) - W(D_{SK} + D_{oil})x_{2}(t)}{\frac{M_{p}}{3} + M_{SK}W} + \frac{-(K + K_{x} + K_{SK}W)x_{1}(t)}{\frac{M_{p}}{3} + M_{SK}W} + (20)$$
$$\frac{3K_{x}bV_{in}(t) + (-1)^{q}a}{(21)}$$

$$\frac{K_x bV_{in}(t) + (-1)^q a}{\frac{M_p}{3} + M_{SK} W}$$
(21)

where q = 1, 2.. The system represented in equations (19) and (21) can be represented as follows:

$$\begin{bmatrix} \dot{x}_1(t) \\ \dot{x}_2(t) \end{bmatrix} = \mathbf{f}(\mathbf{x}(t), H(y(t), V_{in}(t))) + \mathbf{B}V_{in}(t), \quad (22)$$

where it is assumed that  $V_{in}(t) = V_z(t)$ ,

$$\mathbf{f}(\mathbf{x}(t), H(y(t), V_{in}(t))) = \begin{bmatrix} x_2(t) \\ \frac{-Dx_2(t) - (K + K_x + K_{SK}W)x_1(t)}{\frac{M_p}{3} + M_{SK}W} \end{bmatrix}, \quad (23)$$

and  $\mathbf{B} = \begin{bmatrix} 0 \\ \frac{3K_x b + (-1)^q a}{\frac{M_p}{3} + M_{SK} W} \end{bmatrix}$ . The following PID controller is defined:

$$K(t) = \mathbf{G} \big( \mathbf{x}_d(t) - \mathbf{x}(t) \big), \tag{24}$$

where  $\mathbf{G} = \begin{bmatrix} P_i & D_i \end{bmatrix}$ , and  $\mathbf{x}_d(t)$  represents the vector of the desired piezo trajectories. Equation (24) becomes as follows:

$$K_{i}(t) = \begin{bmatrix} P_{i} & D_{i} \end{bmatrix} \begin{bmatrix} x_{1d}(t) - x_{1}(t) \\ x_{2d}(t) - x_{2}(t) \end{bmatrix} + I_{i} \int (x_{1d}(t) - x_{1}(t))dt, \qquad (25)$$

thus

$$K_{i}(t) = P_{i} (x_{1d}(t) - x_{1}(t)) + D_{i} (x_{2d}(t) - x_{2}(t)) + I_{i} \int (x_{1d}(t) - x_{1}(t)) dt, \quad (26)$$

 $P_i$  and  $D_i$  are internal P and internal D parameters of the PID controller. If the following Lyapunov function is defined:

$$V(K_i) = \frac{K_i^2(t)}{2},$$
 (27)

then it follows that:

$$\dot{V}(K_i) = K_i(t)\dot{K}_i(t).$$
(28)

In order to find the stability of the solution s(t) = 0, it is possible to choose the following function:

$$\dot{V}(K_i) = -\eta(t)K_i^2(t),$$
 (29)

with  $\eta > 0$ . Comparing (28) with (29), the following relationship is obtained:

$$K_i(t)K_i(t) = -\eta K_i^2(t),$$
 (30)

and finally

$$K_i(t)(K_i(t) + \eta K_i(t)) = 0.$$
(31)

The no-trivial solution follows from the condition

$$K_i(t) + \eta K_i(t) = 0.$$
 (32)

From (24) it follows:

$$\dot{K}_{i}(t) = \mathbf{G}(\dot{\mathbf{x}}_{d}(t) - \dot{\mathbf{x}}(t)) + I_{i}(x_{1d}(t) - x_{1}(t)) = \mathbf{G}\dot{\mathbf{x}}_{d}(t) - \mathbf{G}\dot{\mathbf{x}}(t) + I_{i}(x_{1d}(t) - x_{1}(t)).$$
(33)

The main idea is to find a  $u_{eq}(t)$ , an equivalent input, and after that a  $V_{in}(t)$ , such that  $\dot{\mathbf{x}}(t) = \dot{\mathbf{x}}_d(t)$ . For that, from (22) it follows that:

$$\dot{\mathbf{x}}(t) = \dot{\mathbf{x}}_d(t) = \mathbf{f}(x_d(t), H) + \mathbf{B}V_{in}(t), \qquad (34)$$

and from (33) the following relationship is obtained:

$$\dot{K}_{i}(t) = \mathbf{G}\dot{\mathbf{x}}_{d}(t) - \mathbf{G}\mathbf{f}(x_{d}(t), H) - \mathbf{G}\mathbf{B}V_{in}(t) = \\ \mathbf{G}\mathbf{B}(u_{eq}(t) - V_{in}(t)) + I_{i}(x_{1d}(t) - x_{1}(t)), \quad (35)$$

where  $u_{eq}(t)$  is the equivalent input which, in our case, assumes the following expression:

$$u_{eq}(t) = \left(\mathbf{GB}\right)^{-1} \mathbf{G}\left(\dot{\mathbf{x}}_d(t) - \mathbf{f}(x_d(t), H)\right).$$
(36)

After inserting (35) in (32) the following relationship is obtained:

$$\mathbf{GB}\big(u_{eq}(t) - V_{in}(t)\big) + \eta K_i(t) = 0, \qquad (37)$$

and in particular

$$V_{in}(t) = u_{eq}(t) + \left(\mathbf{GB}\right)^{-1} \eta K_i(t).$$
(38)

Normally, it is a difficult job to calculate  $u_{eq}(t)$ . If equation (35) is rewritten in a discrete form using Euler approximation, then it follows:

$$\frac{K_i((k+1)T_s) - K_i(kT_s)}{T_s} = \mathbf{GB} \left( u_{eq}(kT_s) - V_{in}(kT_s) \right).$$
(39)

If equation (38) is also rewritten in a discrete form, then:

$$V_{in}(kT_s) = u_{eq}(kT_s) + \left(\mathbf{GB}\right)^{-1} \eta K_i(kT_s).$$
(40)

Equation (39) can be also rewritten as:

$$u_{eq}(kT_s) = V_{in}(kT_s) + (\mathbf{GB})^{-1} \frac{K_i((k+1)T_s) - K_i(kT_s)}{T_s}.$$
(41)

Equation (41) can be estimated to one-step backward in the following way:

$$u_{eq}((k-1)T_s) = V_{in}((k-1)T_s) + (\mathbf{GB})^{-1} \frac{K_i(kT_s) - K_i((k-1)T_s)}{T_s}.$$
 (42)

Because of function  $u_{eq}(t)$  being a continuous one, we can write:

$$u_{eq}(kT_s) \approx u_{eq}((k-1)T_s).$$
(43)

Considering equation (43), then equation (42) becomes:

$$u_{eq}(kT_s) = V_{in}((k-1)T_s) + (\mathbf{GB})^{-1} \frac{K_i(kT_s) - K_i((k-1)T_s)}{T_c}.$$
 (44)

Inserting (44) into (40):

$$V_{in}(kT_s) = V_{in}((k-1)T_s) + (\mathbf{GB})^{-1} \left(\eta K_i(kT_s) + \frac{K_i(kT_s) - K_i((k-1)T_s)}{T_s}\right), \quad (45)$$

and finally:

$$V_{in}(kT_s) = V_{in}((k-1)T_s) + (\mathbf{GB}T_s)^{-1} \left(\eta T_s K_i(kT_s) + K_i(kT_s) - K_i((k-1)T_s)\right).$$
(46)

Through matrix B the control law is a switching one.

#### B. A switching external PID controller

If the following external PID is defined:

$$s_e(t) = \begin{bmatrix} P_e & D_e \end{bmatrix} \begin{bmatrix} x_{1vd}(t) - x_{1v}(t) \\ x_{2vd}(t) - x_{2v}(t) \end{bmatrix} + I_e \int (x_{1vd}(t) - x_{1v}(t))dt, \qquad (47)$$

where  $\mathbf{x}_{vd}(t)$  represents the vector of the desired valve trajectories. Then after similar calculation the following PID structure is calculated:

$$V_{in}(kT_s) = V_{in}((k-1)T_s) + (\mathbf{LH}T_s)^{-1} \left(\eta T_s K_i(kT_s) + K_e(kT_s) - K_e((k-1)T_s)\right).$$
(48)

The following notation is used:  $\mathbf{L} = \begin{bmatrix} P_e & D_e \end{bmatrix}$  and

$$\mathbf{H} = \begin{bmatrix} 0\\ -F(x_{SK_d}(t)y_d(k))/M_v \end{bmatrix},\tag{49}$$

where  $c_{x_{SK_d}(k)} \in \mathbb{R}$  and  $c_{y_d(k)} \in \mathbb{R}$  are the constants which come from the linearization of  $F(x_{SK}(t), y(t))$  at each t-time.

#### IV. SIMULATION RESULTS

The model described in II is considered together with the control law of Eq. (46). In Fig. 9 final results which describe the tracking of a desired position of an exhaust valve with 8000 rpm are presented. In Fig. 10 final results which describe the tracking of a desired velocity of an exhaust valve with 8000 rpm are considered. In Fig. 11 the force acting directly on the valve is shown. The force at the opening time has a peak value equal to 700 N circa and it is reduced to a few Newtons acting on the servo piston part thanks to the decoupling structure of the hybrid actuator. The model obtained with such a disturbance is an exponent function of the position of the valve. The profiles of pressure  $p_A(t)$  and  $p_B(t)$  of the oil chambers of Fig. 7 are shown in Fig. 12.



Fig. 9. Desired and obtained valve positions corresponding to 8000 rpm



Fig. 10. Desired and obtained valve velocity (8000 rpm)

#### V. CONCLUSIONS

#### A. Conclusions

A model of a hybrid actuator is considered in this contribution together with its control. The proposed control strategy consists of an internal PID and an external one to be applied in an engine application. Simulation results consist of tracking of exhaust valve trajectories are shown.

#### REFERENCES

- Po-Kwang Chang Jium-Ming Lin. Eliminating hysteresis effect of force actuator in a spm. WSEAS TRANSACTIONS on CIRCUITS and SYSTEMS, 11(11):351–350, 2010.
- [2] H.J.M.T.A. Adriaens, W.L. de Koning, and R. Banning. Modeling piezoelectric actuators. *IEEE/ASME Transactions on Mechatronics*, 5(4):331– 341, 2000.
- [3] Y.-C. Yu and M.-K. Lee. A dynamic nonlinearity model for a piezoactuated positioning system. In *Proceedings of the 2005 IEEE International Conference on Mechatronics, ICM 2005*, pages 28–33, Taipei, 10<sup>th</sup>-12<sup>th</sup> July 2005.
- [4] H. Murrenhoff. Servohydraulik. Shaker Verlag, Aachen, 2002.



Fig. 11. Force of the internal combustion considering 8000 rpm



Fig. 12. Pressure  $p_A(t)$  and  $p_B(t)$  of the oil chambers considering 8000 rpm

- [5] J. Lee, D. Lee, and S. Won. Precise tracking control of piezo actuator using sliding mode control with feedforward compensation. In *Proceedings* of SICE Annual Conference 2010, pages 1244–1249, Taipei, 18<sup>th</sup>-21<sup>st</sup> August 2010.
- [6] M.Y Ali. Experimental set up verification of servo dc motor position control based on integral sliding mode approach. WSEAS TRANSACTIONS on SYSTEMS and CONTROL, 7(3):87–96, 2013.

# Arachnid-inspired Kinesthesia for Legged Robots

Nicholas Harvey School of Mechanical and Industrial Engineering University of Johannesburg Johannesburg

Abstract—In this paper we investigate the use of a biologically inspired kinesthesia organ from the arthropod family as a way to allow estimation of the leg positions of a mobile legged robot. Such an extra sense would serve the same purpose it does in biological equivalents. The *lyriform organ* of the arachnids is used as the particular model for such a biologically inspired organ (sensillum). A control design and finite element analysis of the resulting structure shows the ability of decoding the leg position from data sensed from sensilla integrated into the surface of a robot's main structure. A method for constructing such a robotic equivalent to a lyriform organ is demonstrated using a 3D printing technique.

#### I. INTRODUCTION

The ability to sense and evaluate an environment is of utmost importance to all animals, from the least developed to the most developed. Biological structures have evolved that enable creatures to obtain information about themselves and their environment, using a multitude of different principles (for example visible light sensing and audio echolocation to name but two vastly different ones). Knowledge of their surroundings enables organisms to navigate effectively, but the use of external data is only one aspect of successful navigation [1], [2]. In order to traverse an environment under varying environmental conditions, a creature must also have knowledge of where its limbs are at any given time [3].

Information about limb position and movement is a sense known as *kinesthesia*, or alternatively, *proprioception* [2]. Knowledge about the position of a limb allows an organism to manipulate the appendage to reach the desired position required for navigation [4]. Many structures exist in nature that provide the brain with kinestheic information, such as the vestibular system or muscle spindles [2], [5]. Not only does kinesthesia (in combination with environmental data) allow for determination of limb position relative to surroundings, but also the relative position of limbs themselves [2]. This information is useful because it prevents limbs from interfering with each other or the body itself, even in the absence of other feedback such as visual perception.

The development of robotics has, and is increasingly, focused on mimicking biological systems [1]. Analogues for a number of natural solutions to movement and sensing have been developed, such as artificial muscles [6]. Traditionally however, information about limb position in robots has been derived from actuator feedback (such as a servo reporting its rotational position, or a piston providing information about its extension), which is then interpreted to determine where a limb is [1]. Although this has not affected the ability of

André Leon Nel School of Mechanical and Industrial Engineering University of Johannesburg Johannesburg Email: andren@uj.ac.za

robots to perform high-accuracy functions (industrial robots, for example [7]), more complicated robotic systems, such as legged robots, could benefit from better and additional limb position estimation mechanisms.

#### II. KINESTHESIA IN BIOLOGY

In any complex organism, there are numerous sensory receptors located throughout the body, that may be activated by various causes [8]. Input to a biological sensor may be exteroceptive (originating from the environment) or interoceptive (originating from within the organism itself) in nature. A special sub-category of interoceptive information is proprioceptive input information about the position of body parts relative to the body itself [8]. In the literature, some scholars reserve these terms (interoception and proprioception) for information regarding the internal state (e.g. temperature) and information regarding limb position respectively. Still other scholars make use of the term exproprioception as the term describing the evaluation of limb position relative to the body and the environment. Finally, the terms proprioception and kinesthesia are used interchangeably by some [4], and by others to refer exclusively to position and motion of limbs respectively.

Spiders (the arachnid family) are unique amongst arthropods in that they have the most extensive array of biological sensors distributed across their exoskeletons. These sensors consist of groups (arrays of as many as 30) of slit sensilla<sup>1</sup> called "lyriform organs" [9]. An example of a lyriform organ can be seen in Figure 1.



Fig. 1. Scanning electron microscope (SEM) image of spider sensilla [10]

These organs are primarily located in the region of the leg joints, and respond to the movements of adjacent limbs

<sup>&</sup>lt;sup>1</sup>In this paper we reserve *sensillum* for an individual slit and *sensilla* for the complete organ comprised of many slits.

[9]. Each slit sensilla is paired with two sensory cells, and this arrangement forms what amounts to a biological strain gauge. As the exoskeleton is exposed to forces (which are typically compressive, since they are caused by the transfer of force from the ground to the exoskeleton via adjacent limbs), exaggerated strain occurs in the region of the slits (the discontinuity in the exoskeleton acting as a natural stress concentration mechanism) which is then detected by the sensilla's sensory cells. According to Barth, strains as small as  $-10\mu\varepsilon$  are capable of triggering responses from the sensory cells of such arachnid sensilla [10]. Schaber et al. used white light interferometry to investigate the sensitivity of the lyrifom organs in Cupiennius Salei, finding that the sensitivity of a slit was proportional to its length, and in general sensitive to a compression of 0.1% the width of the slit [9]. Hobl et al. presented a paper wherein they modelled lyriform organs using the finite element method (FEM), with various orientations and arrangements of slits [11]. They concluded that the arrangement of the slits can alter the sensitivity of the system to various directions of loading, and are orientated and located at certain sites on the legs of spiders in order to be optimally receptive to the strain information that is available.

Zill and Moran conducted an extensive study into the variety of proprioceptive organs found on the exoskeleton of insects [12]. The Campaniform sensilla (an organ common to numerous insects, such as flies and cockroaches) is a mechanoreceptor that transmits information about strain in the exoskeleton, similar in function to the lyriform organs in spiders discussed previously. The organ consists of a hole in the cuticle of the exoskeleton (usually not circular but rather oval in shape, occasionally being long and narrow), with a bell-shaped cap suspended on a sensory cell within the hole, as in Figure 2.



Fig. 2. Schematic of the Campaniform sensilla in a fly [13]

#### III. EXISTING KINESTHESIA IN ROBOTS

Kinesthesia in robotic systems that do not rely on actuator feedback is a largely unexplored field of interest at the present time [14] and [15]. Traditionally, the motion and position of limbs have been derived using feedback from sensors that are directly attached to the actuators (e.g. rotary encoders on motors). Although this method has been sufficient for the development of robotics historically, there is evidence to suggest that this method may no longer be the optimal approach, at least in certain applications. For example, as actuators modelled after biological muscles become more commonplace, and joints develop greater degrees of freedom, it is inevitable that the mechanics governing the movement of a limb will become more difficult to model mathematically [16].

As the structures and mechanisms in robots, especially in humanoid and other systems modelled after biological solutions, imitate these systems more closely, it is inevitable that new solutions to proprioception will have to be developed. Nakanishi et al. note that it is difficult to directly measure the orientation of multiple degree of freedom joints, such as those found in spherical joints being used in the shoulders and hips of humanoid robots [16]. Further, they explain the difficulty of adapting current sensing technology (gyroscopes, accelerometers, Hall-effect sensors) to these applications since the sensors are prone to drift and calibration errors. The authors propose a solution to proprioception in a humanoid robot with actuators similar in principle to the muscle tendons in humans [16]. Limb motion is controlled by motors, pulleys and wires, and hence the relative displacement of each tendon can be measured from the rotary encoder on the motor. The focus of the research by Nakanishi et al. is not on the development of new proprioceptive sensors, but rather on how to use data that is generated by actuator systems as a basis for posture estimation in complex joints.

French and Wicaksono designed a proprioception system based on the campaniform sensillum of the fly [13]. Noting the limitations of then current microelectronic machining processes, their design was a somewhat simplified version of the biological structure, foregoing the dome-shaped structure for a flattened design. The results obtained appeared promising, as they were able to demonstrate the effect of recess geometry on stress concentration, and hence on the accuracy of such a system. Unfortunately however, they did not use their findings to construct an actual proprioception system, but rather to demonstrate a new method of strain sensing.

Kramer et al. recently developed a novel solution to joint angle proprioception consisting of a polydimethylsiloxane (PDMS) film with an embedded microchannel containing a conductive liquid, as well as a sensing element [17]. As the film is deformed during bending, the cross-section of the microchannel is altered and hence the resistance of the liquid changes. This change is then measured by the sensor and used to determine joint position. The use of an elastomer film allows the sensor to operate without interfering with the motion of the system. The advantage of their design is that it can be easily adapted to current robotic systems, to provide actuatorindependent information regarding local limb orientation.

Jaax et al. developed a system intended to mimic the muscle spindles found in mammalian tissue, which act as a combined actuator/sensor [5]. As a basis for their design they identified the elements that represent the core functionality of the biological sensor, and how these elements could be approximated by an artificial sensor. Their solution consisted of a combined actuator/transducer. Transduction of the actuator movement is achieved by a set of strain-gaged cantilevers, mounted perpendicular to the axis of actuation. Experimental testing of the sensor/actuator system when subjected to a sinusoidal displacement produced results in line with the performance desired.

Kang et al. recently developed an extremely sensitive

sensor modelled after sensory slit organs found in spiders, the lyriform organs discussed previously [18]. The authors had the goal of creating a multifunctional sensor that was highly sensitive, flexible and durable. Using the mechanics and principles of the lyriform organs as a basis from which to develop an artificial solution, a sensor was created by depositing a 20nm platinum layer on polyurethane acrylate. Cracks were then induced in the platinum by bending the strip to various curvatures, controlling the density and direction of cracks formed. Sensor output is generated by measuring the resistance of the platinum strip, which changes depending on how the structure is formed resulting in various crack interactions. Depending on whether the strip is extending, contracting or being twisted, the resistance will correspondingly change (as crack faces are pulled apart or pushed together).

#### IV. PROPOSAL FOR THE USE OF KINESTHESIA

One of the most natural uses of kinesthesia is the ability to maintain a pose simply by estimation of body members despite a lack of other sensory feedback such as vision or muscle spindle sensing. To some extent the use of the joint angular feedback in legged robots is already a form of robotic kinesthesia - but of a very limited and wholly localized form. Kinesthesia in humans, is characterized by an ability to recognize very accurately the position of limbs relative to the trunk even in the absence of certain neural pathways to the brain related to the primary motion planning structures in the neural system - whether this deficit is due to injury or disease is not relevant. This facility is certainly able to reduce the direct load on the brain in terms of managing the effort to "remember" limb positions to enable quick reflex based actions. It is expected that adding such an overlay system to a legged robotic structure would result in the following benefits:

- reduced computational load on the central processing unit (CPU) of the robot,
- reduced effort in managing sensor input to the CPU,
- the ability to prevent leg-leg or leg-body interactions due to leg motion as a result of secondary information being available to the robot, and
- the realization of an immanent self-awareness of the total body configuration.

In the final instance it would be useful to incorporate all the information from the various sensilla distributed around the robotic skeleton into a single neural structure that would create the articifical equivalent of a "body sense" facility. Even if not integrated using a neural structure it is posited that by using standard estimation modeling this could be achieved but at the cost of a certain CPU load.

A subsidiary, but equally relevant, issue is the ability to ensure robot survivability after the loss of certain sensor / actuators by being able to estimate body orientations from the sensilla that are essentially passive sensing devices and that would be less prone to failure than active sensors such as tachometers or joint angle sensors.

What we desire therefore is an ability to estimate the overall leg and body position data, including joint data from sensors that are firstly passive and therefore less likely to fail and secondly are not related to the motion planning or performance function of the robot, so that reinforcement of body position and motion can take place. Consider that the load attachment point in Figure 10 is indicative of the motor and leg attachment in a physical robot design.

1) The Proposal: To develop a structurally equivalent sensilla form that mimics (however roughly) a biological system that allows for the recovery of global leg and body position data from the local strain information in such a way that multiple sensilla reinforce each other's inputs to an overall sensing processor that acts independently of the actuator based robotic motion planning and execution structure.

#### V. EXPERIMENTAL PROCESS

The aforegoing discussion requires at least some form of design, implementation and testing cycle to provide experimental proof of the concept for the proposal. The next sections present the three initial experimental steps that have been completed as well as the results achieved. In summary the methodology was:

- Develop a realistic mechanical design that mimics the biologically occurring lyriform organ.
- Produce a FEM analysis to determine the response of such a mechanical sensilla.
- Analyse the FEM outputs for various orientations of the local leg.
- Estimate the ability to determine the load condition by only observing the strain from the sensilla's individual slit outputs.

The parameters of the actual lyriform organ that are to be varied for analysis purposes (for development of a practical rather than a control design) can be listed as: slit orientation relative to load direction; slit length; slit spacing; and slit grouping.

#### A. System Design

Using the practical variations estimated from SEM data we can posit the following as base information for the mechanical design process.

- A control surface near the sensilla that features curvature between the point of leg attachment and the body
- The relative slit orientation to the point of attachment
- The slit width to exoskeleton thickness (forms the  $\pi_1$  dimensionless group<sup>2</sup>)
- Slit length related to the slit width (forms the  $\pi_2$  dimensionless group)
- Slit spacing related to the slit width (forms the  $\pi_3$  dimensionless group)
- Local slit groupings relative to each other

<sup>&</sup>lt;sup>2</sup>Pi dimensionless groups relate to the use of the Buckingham Pi Theorem from Dimensional Analysis[19].

Before concept development began, it was decided that 3D printing would be used to create the physical objects for testing purposes. This decision can be motivated as follows:

- 3D provides new opportunities for engineers and product development, and provides manufacturing capabilities that would not be obtainable otherwise.
- An UP! Plus 2 3D Printer was readily available during the design process.
- Intricate geometry could be created more easily than would have been possible using conventional machining processes.
- Plastic provides a good analogue to an exoskeleton, since it is lightweight but strong enough for robotic structural applications.

The decision to use an UP! Plus 2 for manufacturing of the sensilla test element imposed certain restrictions on the design process. These restrictions are listed below:

- Although nearly any geometry can be created, it is necessary to take into account the nature of an extrusion layer manufacturing process, i.e. that thin slits (relative to the layer thickness) cannot be printed in a vertical orientation.
- There is a limit on how thin the slits may be, which is dictated by the nozzle diameter and layer thicknesses used, it was determined that the minimal slit width was about 0.5mm.
- Since parts are printed on a raft and base, there is a small amount of fusing between the actual part and the support structures.
- Although there are an unlimited number of theoretical geometries that are achievable with the UP!, environmental conditions may have a significant impact and hence, the following design constraints were maintained throughout the concept generation phase:
  - slit width would be limited to greater than 0.5mm affects all the  $\pi$  groups, and
  - skeletal thickness would always be greater than 1.0mm (this translates to a slit depth of 1.0mm and affects the  $\pi_1$  group).

Given the restrictions of the 3D printer we have access to the design was limited to having 1.0mm minimum wall / slit thicknesses. This implies that all three the  $\pi$  dimensionless groups for the manufactured product at this stage are violated and that the control design is relatively stiff compared to the actual lyriform organs - and especially stiff compared to any practical lyriform organ that would be added to a real legged robot such as in [20].

#### B. The SEM and Robotic model

As can be seen in Figure 1 it is clear that sensilla vary substantially even on a single specimen. The SEM image does however indicate that most of the sensilla are located in areas adjacent to places of large curvature in the exoskeleton.

As a first version of a sensilla equivalent structure we considered a design as shown in Figure 3. A localized section

of a possible robotic structure where the actuator for a leg is attached is used. By inspection of the details of the sensilla visible in Figure 1 that the sensilla are located in areas where the exoskeleton is curved. The reason for the location of the sensilla around areas of curvature may only become clear on analysis of the results finally achieved by the FEM analysis. In the figure (Figure 3) the design shows the resemblance between the artificially created slit based openings in the robotic structure and the naturally occurring sensilla of the exoskeleton.



Fig. 3. Initial base design of robotic sensilla (Control Design)

#### C. FEM Analysis

Figure 3 shows the design of a possible first generation of a robot sensilla on a simple shell that represents a spider exoskeleton which can be used as the basis of a first FEM analysis as shown in Figure 4. This particular design and FEM analysis will function as the control for further and later design refinements. The FEM analysis as shown has been tested for convergence and for stability. The natural stiffness ratio of a spider exoskeleton (as described by the various  $\pi$  groups) could not be reasonably achieved with standard FEM analysis but allowance for the too stiff structure has been made in analysing the results achieved. A practical printed version of the Control Design is shown in Figure 5.

The FEM analysis shows results that could support the ability to sense leg / joint states based on the use of strain sensing such as the platinum cracked surface sensors of Kang et al. [18] with the present structural stiffness of the openings. From the initial stress / strain analysis (even though it is at present a simple linear analysis) it is obvious that the outputs shown in Figure 6 are adequate for preliminary design reviews.



Fig. 4. Structural modeling of the initial base design

Figure 6 reveals the typical deformation patterns due to a shear load at the exoskeleton. This would represent the deformation to be expected from a leg that is in the air (at



Fig. 5. Practical Control Design



Fig. 6. FEM outputs for the Control Design - shear load applied

a particular angle with respect to the vertical) with no ground contact. In Figure 7 we can see that deformations resulting from a static moment load at the exoskeleton that is typical of a support load for a structure - indicative that the leg is in contact with the ground and resisting some force. Of course since this FEM analysis is strictly linear it is possible to recover intermediate values and combinations using the superposition principle.

Figure 8 and Figure 9 display FEM values for the deformations / displacements to be expected from the practical sensilla - and this is a target to be achieved by the sensors to be incorporated into a future practical design.

#### D. Experiments

The way that the practicality of the Control Design was checked was to substantiate the FEM analysis by loading an actual 3D printed module and using the experimental setup shown in Figure 10. In this way the ability to recover load direction and type from the sensilla motion can be checked. Initial results show that this is being done quite adequately to allow for recovery of both shear and moment based loads at



Fig. 7. FEM outputs for the Control Design - moment applied



Fig. 8. Deformations due to shear in the Z direction

points close to the leg-body joint by the sensilla displacements.

#### VI. DISCUSSION

The Control Design presented has demonstrated that it is adequate to test the proposed solution - even though the particular deformation results for the given shell are too small by an order of magnitude at present for all but analysis purposes. Developing the actual sensors and neural acrchitecture are the necessary next steps in this research.

- It is possible to create a simple slit formed sensilla that replicates a lyriform organ.
- A possible design makes it clear that simple strain sensing - that could be seen as a parallel for direct neural sensing - is possible in a manner that returns outputs that are not simply localized



Fig. 9. Deformations in the X direction due to a moment applied



Fig. 10. Experimental setup for testing of Control Design robotic sensilla (1) - Support bracket; (2) - Test piece with shear attachment modification; (3) - Moment load attachment

- An actual 3D printed version of a part of such a legged robotic exoskeleton including the sensilla proves the practicality of the proposal.
- A detailed FEM analysis of the 3D printed form makes it clear that results are not strictly localized.

#### VII. CONCLUSION

By starting from the assumption that kinesthesia and its implementation via sensilla in a legged insect world has an evolutionary value we propose that to further decrease the gap between natural and artificial robots the evolution of a kinesthetic mechanism in legged robots should be attempted. In this paper we recognize the structure of biological sensilla that allows and use this as the basis for an introductory design of such a kinesthetic mechanism. The control design has been analysed using a FEM analysis so that the potential efficacy of such a sensor can be evaluated. Initial results appear to indicate that the existing structure fulfills all the requirements of such a biologically based sensilla in the locomotion of an insect.

#### REFERENCES

[1] G. A. Bekey. Autonomous Robots: From Biological Inspiration to Implementation and Control, chapter Control and Regulation in Biological Systems, pages 7–43. MIT Press, 2005.

- [2] P. Corke and J. Dias. An introduction to inertial and visual sensing. *The International Journal of Robotics Research*, (6):519–535, 2006.
- [3] B. L. Riemann and S. M. Lephart. The sensorimotor system, part ii: The role of proprioception in motor control and functional joint stability. *Journal of Athletic Training*, (37):80–84, 2002.
- [4] L.A. Jones. *Human and Machine Haptics*, chapter Kinesthetic Sensing. MIT Press, 2000.
- [5] K.N. Jaax, P.H. Marbot, and B. Hannaford. Development of a biometric position sensor for robotic kinesthesia. *Proceedings of the 2000 IEEE International Conference* on Intelligent Robots and Systems (IROS), 2000.
- [6] S.S. Ge and F.L. Lewis. Autonomous Mobile Robots, chapter Sensors and Sensor Fusion, pages 5–41. Taylor & Francis, 2006.
- [7] M.P. Groover. Fundamentals of Modern Manufacturing, chapter Automation Technologies for Manufacturing Systems, pages 887–917. Wiley, 2010.
- [8] G. A. Bekey. Autonomous Robots: From Biological Inspiration to Implementation and Control, chapter Ascending Sensory Pathways. MIT Press, 2005.
- [9] C. F. Schaber S. N. Gorb and F. G. Barth. Force transformation in spider strain sensors: white light interferometry. *Journal of the Royal Society Interface*, pages 1254–1264, 2012.
- [10] F. G. Barth. Frontiers in Sensing: From Biology to Engineering, chapter Spider strain detection, pages 251– 273. Springer, 2012.
- [11] B. Hobl, H.J. Bohm, F.G. Rammerstorfer, and F.G. Barth. Finite element modeling of arachnid slit sensilla - i. the mechanical significance of different slit arrays. *Journal* of Computational Physiology, pages 445–459, 2006.
- [12] S. N. Zill and D. T. Moran. The exoskeleton and insect proprioception. i responses of tibal campaniform sensilla to external and muscle-generated forces in the american cockroach, periplaneta americana. *Journal of Experimental Biology*, pages 1–24, 1981.
- [13] P. French and D.H. Wicaksono. Biologicallyinspired mechanical sensors. Available from http://ei.et.tudelft.nl/assignments/masterprojects/2007/french/, 2015.
- [14] F. Tang. Cs599: Robotics winter 2013. Available from http://www.cpp.edu/ ftang/courses/CS599/notes/, 2015.
- [15] R.R. Murphy. Introduction to AI Robotics, chapter Common Sensing Techniques for Reactive Robots, pages 207–209. MIT Press, 2000.
- [16] I. Mizuuchi Y. Nakanishi, K. Hongo and M. Inaba. Joint proprioception acquisition strategy based on jointsmuscles topological maps for muscoskeletal humanoids. *Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- [17] R. Sahai R. K. Kramer, C. Majidi and R. J. Wood. Soft curvature sensors for joint angle proprioception. *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.
- [18] D. Kang, P.V. Pikhitsa, Y.W. Choi, C. Lee, S.S. Shin, L. Piao, B. Park, K.-Y. Suh, Kim T., and M. Choi. Ultrasensitive mechanical crack-based sensor inspired by the spider sensory system. *Nature*, 516(7530):222–226,

December 2014.

- [19] J. Bertrand. Sur l'homognit dans les formules de
- [19] J. Bertrand. Sur Fnomogint dans its formules de physique. *Comptes rendus*, 86(15):916920, 1878.
  [20] S.T. Marais, F. du Plessis, and A.L. Nel. Architecture for a hexapod robot. *Proceedings of the 2015 RobMech* Conference (Poster session), 2015.

## A straightforward method for calculating the voicing cut-off frequency for streaming HNM TTS

J.A.Louw

Human Language Technologies Research Group, Meraka Institute, CSIR, Pretoria, South Africa jalouw@csir.co.za

*Abstract*— The Harmonic plus Noise Model vocoder produces natural text-to-speech synthesis without some of the artifacts encountered in other vocoders. However, in order to achieve this naturalness one needs to determine a voicing cut-off frequency for each frame of speech being synthesized. This has proven to be a challenge and there are many methods and implementations, all with certain trade-offs. We present here a straightforward method, based on cepstral energy, that can also be used in streaming HNM TTS synthesis.

#### I. INTRODUCTION

In a text-to-speech (TTS) system the *vocoder* (also described as the speech generation module) is generally the last functional block and generates the synthetic waveform from the specifications created in all the previous blocks or steps. This can be as simple as concatenating pre-recorded wave forms, or an involved process where the speech wave from is generated from a sequence of spectral parameters or features. Generally the naturalness of a speech synthesis system, in terms of synthetic artifacts (buzzing, clipping, or other synthetic sources of noise), can be ascribed to the vocoder.

A popular vocoder in TTS is the *Harmonic plus Noise Model* (HNM) [1]. The HNM is a complete parametric model of a frame of speech, where the speech spectrum is decomposed into two parts: a harmonic component and a noise component. The vibration of the vocal folds during speech production are modeled by the harmonic part, while the modulation of the airflow during speech production is modeled by the noise part.

Fig. 1 (adapted from [2]) shows an example of a frame of speech where the separation between the harmonic and noise components of the speech as modeled by HNM are clear and the harmonic spectral peaks are marked with an asterisk (\*).



Fig. 1. An example of the power spectrum of a frame of speech with a clear separation between the harmonic and noise parts [2].

The point where the spectral separation occurs is called the *Voicing Cut-Off Frequency* (VCO) or the *Maximum Voiced Frequency* (MVF). In this paper we will refer to it as the voicing cut-off frequency. The VCO is a time varying parameter and the location thereof is specific to the phone being reproduced, co-articulation effects, the within phone position, the macro intonation, etc. In practice this separation is not always as clear as depicted in Fig. 1, and one can find multiple sections of a spectrum of a frame of speech with harmonic and noise characteristics, such as depicted in Fig. 2 (adapted from [3] with a calculated VCO of 3.6 kHz). Therefore, the definition of the VCO is an ill posed problem, and its estimation a heuristic process [3].



Fig. 2. An example of the power spectrum of a frame of speech where the separation between the harmonic and noise parts is not clearly distinguishable [3].

Accurate estimation of the VCO and the evolution of this point in the speech spectrum is important for HNM-based vocoders, as it determines the part of the spectrum that is to be reproduced with a harmonic model and the part that is to be reproduced with a noise model. This paper explores different techniques for the determination of the VCO, and then puts fourth a straightforward method that can also be used in streaming HNM TTS.

The paper is organized as follows: In the next section we give background and related work, then we present our experiment, the results and finally a conclusion.

#### II. BACKGROUND AND RELATED WORK

In HNM the lower band, or the harmonic part is modeled as a sum of the harmonics of the fundamental frequency of the frame of speech in question [4] as

$$s_h(t) = \sum_{k=-L(t)}^{L(t)} A_k(t) e^{jk\omega_0(t)t}$$
(1)

where L(t) denotes the number of harmonics included in the harmonic part,  $\omega_0(t)$  denotes the fundamental frequency and  $A_k(t)$  the time varying harmonic amplitudes. The number of harmonics, L(t), is dependent on the fundamental frequency,  $\omega_0(t)$ , and the voicing cut-off frequency, VCO(t)

$$L(t) = \frac{VCO(t) \cdot 2\pi}{\omega_0(t)} \tag{2}$$

The upper band of HNM, or noise part, can be described in the frequency domain as a time varying autoregressive (AR) model [4],  $h(\tau, t)$ , with its time domain structure imposed by a parametric envelope, e(t),

$$s_n(t) = e(t)[h(\tau, t) \star b(t)] \tag{3}$$

where  $\star$  denotes the convolution of  $h(\tau, t)$  and b(t)and b(t) is white Gaussian noise [4]. The complete HNM synthetic signal is then the sum of (1) and (3)

$$\hat{s}(t) = s_h(t) + s_n(t) \tag{4}$$

The influence of the VCO on the HNM synthetic signal can be seen from (2). When the estimated VCO is higher than the "true" VCO, a larger portion of the synthetic HNM signal is reproduced in the lower harmonic band, while if the estimated VCO is lower than the "true" VCO, a larger portion of the HNM signal is reproduced in the upper noise band. This error effect has a pronounced influence on the perceived naturalness of the generated synthetic speech signal, in that a too low estimate for VCO produces "hissed" voiced sounds, while a too high estimate produces "buzzy" unvoiced sounds.

Previous work on VCO and band splitting can be grouped into the following methods [3]:

**Analysis-by-Synthesis (AbS)** This group of methods or algorithms determine a goodness-of-fit of the speech being analyzed when compared to a specific model. An example can be found in [5]. According to [3], this method does not take into account the distribution of harmonic and noise energy along the frequency axis, and therefore there is a chance that the VCO can be determined to be in the middle of well defined harmonics or a clearly noise part of the spectrum.

**Spectral domain methods** In these methods the spectrum of the frame of speech in question is inspected and each spectral peak is determined to be either voiced or unvoiced by some heuristic criteria. The point where "harmonicity" disappears is determined to be the VCO. An example of these methods can be found in [4].

**Time domain methods** These methods estimate a measure of periodicity on the time signal of the speech frame in question. An example of this method can be found in [6]. [3] states that a drawback of this method is the objective function which often does not have a clear maximum.

In [7] a simple VCO determination was explored, where the VCO for a specific frame of speech was made energy dependent by means of a heuristic linear relationship between the frame's energy and that of the target utterance to be synthesized:

$$v^{(k)} = v^{(max)} \cdot \frac{c_0^{(k)} - c_0^{(min)}}{c_0^{(max)} - c_0^{(min)}} \tag{5}$$

where  $c_0^{(k)}$  is the 0th cepstral coefficient at frame k,  $c_0^{(max)}$  and  $c_0^{(min)}$  are the maximum and minimum values of the 0th cepstral coefficient over all frames of the target utterance.  $v^{(max)}$  is the maximum expected VCO for the specific speaker. They noted that the main disadvantage of this method is its incompatibility with real-time speech waveform generation, because the  $c_0$  range of the whole utterance needs to be known in advance.

This method was compared to an explicit VCO analysis based on a sinusoidal likeness measure (spectral domain method). An error function estimates the error of assuming a specific value for the VCO based on an empirically-adjusted function. Once all VCO candidates have been selected for all frames, the final decision is made by minimizing a cost function through a dynamic programming search.

The results of a perceptual evaluation as given by [7] is repeated in Table I

 TABLE I

 Subjective relative preference for two VCO related

 strategies (adapted from [7]).

Method \Voice	Female	Male
Analytic	30%	29%
No preference	53%	49%
$c_0$	17%	22%

#### **III. EXPERIMENT**

Calculating the chi-squared test on a contingency table [8] from the given data in Table I returns a P-value of 0.813563 at a significance level at p < 0.01. Therefore any differences in cell frequencies could be explained by chance.

In [7] it is stated that the difference in the test samples were not easily discerned by the test subjects, and that the main advantage of the more complicated sinusoidal likeness measure is due to the fact that it can be used during real-time synthesis.

This leads us to believe that the simple cepstral energy method in (5) is sufficient for the calculation of the VCO. In the experiment we compare three different methods of using the simple cepstral VCO method, of which two can be used during real-time synthesis.

The experiment consisted of building three versions of the same TTS voice, each with a different approach towards VCO determination:

- The first voice is built with the cepstral VCO as described in [7] and (5). We label this voice as *Local VCO*.
- The second voice is built with an extra training stream for the HMM-Based Speech Synthesis System (HTS) [9]. This stream consists of VCO data as calculated for each utterance in the training data. Therefore, the parameters  $c_0^{(max)}$  and  $c_0^{(min)}$  are extracted from the training data and an extra information stream is trained, in the same way that training occurs for the  $log(F_0)$ stream for HTS. We label this voice as *Trained VCO*.
- The third voice is built as the first voice, but  $c_0^{(max)}$  and  $c_0^{(min)}$  are calculated on a global scale, in other words from the whole training database, on all utterances:

$$c_{0_{global}}^{(max)} = \frac{\sum_{i=1}^{I} c_{0_{local}}^{max}}{I}$$
(6)

$$c_{0_{global}}^{(min)} = \frac{\sum_{i=1}^{I} c_{0_{local}}^{min}}{I}$$
(7)

 $c_{0_{global}}^{(max)}$  and  $c_{0_{global}}^{(min)}$  are calculated as the arithmetic mean of  $c_{0}^{(max)}$  and  $c_{0}^{(min)}$  of each utterance and *I* is the number of utterances in the training data. Equation (5) then becomes

$$v^{(k)} = v^{(max)} \cdot \frac{c_0^{(k)} - c_{0_{global}}^{(min)}}{c_{0_{global}}^{(max)} - c_{0_{global}}^{(min)}}$$
(8)

This voice is labeled as Global VCO.

For all the voices we set  $v^{max}$  to 4.5 kHz.

#### A. Data

The speech database consists of 5 hours of speech from a professional voice artist, recorded in a studio and sampled at 44.1 kHz. The voice artist is female with a South African English accent. The first two sections of the database are from the Arctic [10] database, sections A and B and consists of 1131 utterances. The following table lists the source material of the database recordings, together with the recording style and the number of utterances.

TABLE II Description of the database.

Source material	Style	# utterances
Arctic	flat	1131
Comma gets a cure	expressive	1
Dates & days	flat	6
Digits	flat	66
General	normal	256
News	read out	100
Novel	expressive	100
Numbers	flat	67
Paragraphs	expressive	209
Addresses	flat	100
Street names	flat	89

#### B. Alignment and HMM Training

The Hidden Markov Model Toolkit (HTK) [11] was used in the forced alignment of the audio to the phonetic transcriptions of utterance in the recorded data (see III-A). These alignments are then used to train speaker-specific triphone acoustic models. The final training phase consists of utterance realignment using the speaker-specific models. The realigned data is then ready for input to the HTS voice training procedure. For the second voice of the experiment an extra stream is defined in order to train on the cepstral  $c_0$ values.

#### **IV. RESULTS**

Informal listening tests suggested that there is no clear preference for one of the three methods and that the differences between the synthesized samples were difficult to perceive. Three speech experts and three nonexperts evaluated a sample of five synthesized utterances of each voice and could not discern between the three voices.

The graphical results given here are for the synthesized utterance "but here amongst ourselves let us speak out". Fig. 4 gives the VCO curves over time for the three voices, with Global VCO corresponding to the case of  $c_{0_{global}}$ , Local VCO to  $c_{0_{local}}$  and Trained VCO to  $c_{0_{stream}}$ . Fig. 5 show the spectrograms of the three trained voices for the utterance "but here amongst ourselves let us speak out".

The *Local VCO* and *Global VCO* methods display a wider dynamic range than the *Trained VCO* method. The *Trained VCO* method has a limited range, which we believe is due to the "over-smoothing" phenomenon in the time-domain caused by the HMM model structure and constraints of the parameter generation algorithm with regards to the static and dynamic features of the modeled variable, as explored in [12].

In Fig. 4 the VCO values of *Global VCO* and *Local VCO* are indistinguishable and seem to track each other. We believe that this can be explained when the distribution of  $c_0$  is taken into account, as depicted in Fig. 3 (adapted from [13]). The distribution of  $c_0$  is bimodal, with clearly distinguishable minima and maxima corresponding to  $c_0^{min}$  and  $c_0^{max}$  respectively. Thus, there is a relatively high probability that  $c_{0global}^{max}$  is close to the value of  $c_{0local}^{max}$  and the same for  $c_{0global}^{min}$  and  $c_{0local}^{min}$ , giving rise to the VCO values which are almost the same (from equations (5), (6) and (7)).

#### **V. CONCLUSIONS**

We showed that a straightforward cepstral energy method can be used to produce high quality voicing cut-off frequencies for Harmonic plus Noise Modeling in textto-speech synthesis. Out of the three methods that were compared the *Global VCO* method is preferred, as this method is also attractive for use in streaming HNM TTS synthesis, due to the fact that the  $c_{0_{global}}$  minimum and maximum values can be determined in an offline mode,



and the run-time computational complexity is low. The  $c_0$  values for the whole utterance is required in advance for the *Local VCO* method, rendering it incompatible with real-time speech waveform generation, while the *Trained VCO* method requires an extra HTS data stream, making it computationally more expensive than the *Global VCO* method and requiring more storage space.

Future work to be explored is the automatic determination of  $v^{max}$  and testing this method on a larger scale with a wider variety of voices.

#### REFERENCES

- J. LaRoche, Y. Stylianou, and E. Moulines, "HNM: A simple, efficient harmonic+ noise model for speech," in *Applications of Signal Processing to Audio and Acoustics*, 1993. Final Program and Paper Summaries., 1993 IEEE Workshop on. IEEE, 1993, pp. 169–172.
- [2] Y. Stylianou, "Harmonic plus Noise Models for Speech, combined with Statistical Methods, for Speech and Speaker Modification," Ph.D. dissertation, 1996.
- [3] K. Hermus, H. Van hamme, and S. Irhimeh, "Estimation of the Voicing Cut-Off Frequency Contour Based on a Cumulative Harmonicity Score," *IEEE SIGNAL PROCESSING LETTERS*, vol. 14, no. 11, p. 820, 2007.
- [4] Y. Stylianou, "Applying the harmonic plus noise model in concatenative speech synthesis," *Speech and Audio Processing, IEEE Transactions on*, vol. 9, no. 1, pp. 21–29, 2001.
- [5] R. McAulay and T. Quatieri, *Speech Coding and Synthesis*. Elsevier, 1995, ch. Sinusoidal coding, pp. 121–173.
- [6] E. K. Kim, W. J. Han, and Y. H. Oh, "A new band-splitting method for two-band speech model," *Signal Processing Letters, IEEE*, vol. 8, no. 12, pp. 317–320, 2001.
- [7] D. Erro, I. Sainz, E. Navas, and I. Hernaez, "Harmonics plus noise model based vocoder for statistical parametric speech synthesis," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 2, pp. 184–194, Apr. 2014.
- [8] B. J. Winer, D. R. Brown, and K. M. Michels, *Statistical principles in experimental design*. McGraw-Hill New York, 1971, vol. 2.
- [9] H. Zen, T. Nose, J. Yamagishi, S. Sako, T. Masuko, A. Black, and K. Tokuda, "The HMM-based Speech Synthesis System (HTS) Version 2," in *Sixth ISCA Workshop on Speech Synthesis SSW6*, Bonn, Germany, August 2006.
- [10] J. Kominek and A. W. Black, "CMU ARCTIC databases for speech synthesis," Language Technologies Institute, Carnegie Mellon University, Tech. Rep. CMU-LTI-03-177, 2003.
- [11] S. J. Young, G. Evermann, M. J. F. Gales, T. Hain, D. Kershaw, G. Moore, J. Odell, D. Ollason, D. Povey, V. Valtchev, and P. C. Woodland, *The HTK Book, version 3.4.* Cambridge, UK: Cambridge University Engineering Department, 2006.

- [12] M. Zhang, J. Tao, H. Jia, and X. Wang, "Improving HMM based speech synthesis by reducing over-smoothing problems," in *Chinese* Spoken Language Processing, 2008. ISCSLP'08. 6th International Symposium on. IEEE, 2008, pp. 1–4.
- [13] J. Du and R.-H. Wang, "Cepstral shape normalization (CSN) for robust speech recognition," in Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on. IEEE, 2008, pp. 4389–4392.



Fig. 4. Spectrograms for utterance with (a)  $c_{0_{local}}$ , (b)  $c_{0_{stream}}$  and (c)  $c_{0_{global}}$  cepstral energy parameters.



Fig. 5. Spectrograms for utterance with (a)  $c_{0_{local}}$ , (b)  $c_{0_{stream}}$  and (c)  $c_{0_{global}}$  cepstral energy parameters.

# Validating a Reconfigurable Assembly System Utilizing Virtual Commissioning

HJ Vermaak

Department of Electrical, Electronics and Computer Engineering, Central University of Technology, Private Bag X20539,Bloemfontein, South Africa, 9300. Email: hvermaak@cut.ac.za

Abstract-South African manufacturing companies today need to be more sophisticated technologically to compete for global markets. The latest trend in automation and manufacturing emerges in the form of reconfigurable systems. The aim of this paper is to show the development of a reconfigurable assembly system and using virtual commissioning to plan, validate and optimize it. To achieve this "DELMIA" software was used to create a virtual simulation environment to verify an assembly cell from such a system as a case study. Simulations were conducted to verify software functions, device movements and operations, and the control software of the system. As a result, it was found that virtual commissioning is an excellent tool for predicting how the system will function, verifying system code early, and rectifying design flaws. This will enable manufacturing companies to be more competitive, ensure increased productivity, save time and ensure them an advantage over their competition.

Keywords— DELMIA, Virtual Commissioning, Digital Manufacturing, Reconfigurable Assembly Systems

#### I. INTRODUCTION

By definition, virtual commissioning is the simulation of a virtual system within a virtual environment without needing to develop a physical system beforehand. By utilizing virtual commissioning, design flaws can be rectified early in the design stage, space reservation can be allocated for the machinery used in the system and controller software verified well in advance, before building the physical system [1]. Virtual commissioning also allows for easy reconfiguration of an existing system, where process, software or hardware changes can be made in the digital model of the system, then analysed to see how these changes influence the system, and then based on the analysis results the physical system can be modified, preventing costly downtime of the physical system. Furthermore, in industry it enables manufacturers to streamline an assembly line, where planning is done more efficiently and through-put can be predicted due to the visualising of the assembly line [2]. In short, virtual commissioning is established when a virtual factory is controlled by a physical programmable logic controller (PLC) to emulate the behaviour of the physical system [3-6].

JA Niemann

Department of Electrical, Electronics and Computer Engineering, Central University of Technology, Private Bag X20539, Bloemfontein, South Africa, 9300. Email: jniemann@cut.ac.za

#### II. PRELIMANARY STEPS TO OBTAIN VIRTUAL COMMISSIONING

To obtain virtual commissioning, some preparations must be done ahead of time. These preparations include that the user must acquire prior knowledge of the simulation software environment under discussion, the devices to be validated along with knowledge of their behaviour and kinematics must be known, and a server [4, 7] to communicate between the control logic and these devices must be installed. Afterwards, the geometry used to represent these devices must be designed, assembled into smart-devices and then connected to the control logic to validate the system behaviour and control software via an execution environment [8].

Subsequently follows how geometry is obtained, how it fit together, how behaviour is assigned to it, and finally how it is used to obtain virtual commissioning and validating a system.

#### A. Product Hierarchy

Any geometry which consists of multiple parts, or other products, is known as a product. A product is the root element of a hierarchy and contains multiple sub-elements to represent the branches of the hierarchy tree. Fig. 1 shows clearly that a collection of parts are grouped together to form an assembly. Here, constraints and tasks can be allocated to the assembly, then be connected to internal device logic to form a smartdevice. In a similar way, multiple smart-devices can be used together to form work cells and ultimately an entire system [9-11].

#### B. Parts, Assemblies, Mechanisms and Tasks

As seen in the previous section, smart-devices consist of geometry and internal logic (IL). The geometry of a smartdevice consists of parts, assemblies and mechanisms (joints). Parts are the most fundamental elements of any geometry, and there are two methods of attaining them, downloading the parts from a vendors' website and alternatively, parts can be



Fig. 1. Hierarchy of a work cell

designed and created by using the suitable computer-aided design (CAD) software [3, 12]. This hierarchy of a work cell is shown in Fig. 1.

In contrast, an assembly is a collection of parts, linked together by means of specified constraints (see Fig. 2). These assemblies represent a mechanical assortment, which contain at least one fixed part and various moving parts. An example is a piston that is driven by a crank. Moreover, the steps taken to construct these assemblies can better be explained by referring to Fig. 2. Firstly, the parts are imported into the environment (shown at number 1 in Fig. 2). Next the parts must be assigned constraints, which include alignment, orientation of parts and surface contact constraints (shown at number 2 and 3 in Fig. 2). After constraint allocation, the geometry can be updated, causing the parts to rearrange into their intended positions (shown at number 4 in Fig. 2). In addition to constraints, these assemblies can be allocated kinematical commands, by specifying the physical limits, direction of movement, speed, and acceleration of the assembly. This is known as creating a joint or mechanism.



Fig. 2. Assembly of parts

Additionally, these mechanisms can be given behaviour, by allocating tasks and operations to it. To differentiate, an operation depicts the movement of a mechanism or several mechanisms when executed, whereas a task executes a series of consecutive operations and functions (like closing a gripper). To allocate a task, the sequence of possible activities are identified and taught to the mechanism. This entails that the mechanism is jogged or moved into place, where after the set of coordinates are recorded in an operation table which will perform the movements sequentially. This will give mechanisms the needed behaviour to be used as smart-devices.

In a similar fashion to joint creation, different smartdevices can be assembled together to form a larger, more complex assembly or system [1]. Fig. 3 shows how separate smart-devices are first imported (shown at number 1), then aligned into position (shown at number 2) and finally attached. These attachments can be seen in the figure at number 3. The smart-devices are attached in a parent-child manner (shown at number 3), which causes smart-devices to move respective to others. The example in the figure shows that the Y-axis (child) must move as if it is fastened to the slider of the X-axis (parent). Likewise, the cylinder (child) moves with the Y-axis (parent) and the gripper tool (child) moves with the cylinder (parent) and so forth [3, 8].

#### C. Control Logic

In order to control device geometry so that it imitates the behaviour of its physical counterpart or cause several devices to cooperate with each other within a digital factory, control logic must be developed. Fig. 4 shows that control elements or logic can be separated into two types namely device logic and



Fig. 3. A completed smart-device



Fig. 4. Device and control logic

control logic, where control logic divides into internal and external control logic.

Firstly, device logic also known as internal logic, assigns unique behaviour to the devices or assemblies created as in the previous section to make it smart-devices. The IL uses the inputs and outputs of each device to control the actions it must perform, which gives each device a distinct behaviour (shown in Fig. 5).

On the contrary, control logic is a standalone supervisory control instance, which can be used to control multiple smartdevices within a digital factory. To clarify, internal control can be seen as a virtual PLC (running inside the environment) connected to devices to emulate the behaviour of a real PLC. With external control on the other hand, the devices are connected to a real system PLC via an OLE for Process Control (OPC) server. The latter is used to validate system PLC code.

#### III. IMPLEMENTATION AND VALIDATION

The steps to obtain virtual commissioning can better be explained by referring to Fig. 6. Firstly all the parts for the system must be obtained by either designing or downloading them. Afterwards, the parts must be assembled into assemblies to build up the devices of the system. The devices can then be given behaviour by defining constraints and allowable movements. At this stage the devices must be verified for accurate mechanical operation by performing joint simulations.



Fig. 5. Control block with internal logic

Next, the internal logic for each device to be validated are developed and then applied to it. This represents the intelligence of each device (internal code of devices). After this step is complete, the geometry can be referred to as smartdevices. Now all the smart-devices can be imported into the environment to complete the entire system to be validated. At this stage the environment is set up for virtual commissioning [8, 13].

To demonstrate how virtual commissioning can aid the validation of a system and its PLC code, the validation procedure is divided into several tests. An optional simulation can be done to verify the operation of the virtual system in the form of a process plan simulation. Operational tasks are assigned to each smart-device to emulate the overall process flow of the system [14]. This method does not validate any control or devices logic, only the functional operation of the virtual system and is not handled in this paper. For test 1 a simulation is set up to validate the virtual version of the system along with device logic. This is done to predict how the virtual system will react to control code. This test is only done to validate the behaviour of the virtual system (device operations and internal logic). In test 2 the virtual system is connected to the physical system PLC via an OPC server. Here simulations are repeatedly done until the operation of the virtual system is satisfactory and thus the physical PLC code verified successfully. This is known as virtual is commissioning. For the last test, the physical system is connected to the verified PLC and operated. Now the operation of the system can be compared to the predictions made through virtual commissioning.



Fig. 6. Steps to virtual commissioning

This will evaluate if virtual commissioning successfully predicted and validated the operation of the physical system [5, 15-17].

#### A. Setup Overview

To achieve this, the setup is as shown in Fig. 7. The virtual version of a device is built and programmed with its behaviour, on a computer with the simulation software environment installed on it. In addition, the simulation environment is connected to an OPC server, which is installed and also runs on the same computer. Furthermore, the server connects to the system PLC via an Ethernet connection, which enables the server to access and communicate with the PLC. To complete the setup, the physical version of the device is built and also connected directly to the main PLC, which enables it to control it.

#### B. Execution Environment and Simulation

In order to set up the simulation, an execution environment must be created. Firstly, the virtual device must be designed, assembled and imported into the environment. Next, the virtual equipment is programmed with the behaviour of the physical equipment. This will transform the virtual device into a smart-device. In addition, joint simulations must repeatedly be done until the operations and movements of the virtual device can imitate its real counterpart based on input/output (IO) signals from a PLC. After the virtual device is perfected, the PLC code to control the system is developed on the actual system PLC. When the PLC code is completed, it is necessary to establish a connection with the OPC server. To achieve this, the IOs of the PLC must be mapped to the virtual device. This can be seen in Fig. 8. Here the IO signals of the PLC are mapped to the virtual equipment exactly the way it should be physically wired to the real equipment.



Fig. 7. System validation overview



Fig. 8. Mapping devices to PLC

When all the connections are made, the execution environment can be compiled and checked for errors and signal strength. If no errors are present and the signal quality of the OPC server is good, the environment can be simulated to validate the PLC code and predict the movements of the system based on the instructions from the PLC. At this stage the PLC code can be modified until the virtual system delivers the desired intended operation of the system. If the virtual environment functions flawlessly with the PLC code, the PLC can be connected directly to the physical equipment. At this stage the validated PLC code runs and controls the physical system in real-time. Now a comparison can be made between the prediction made through virtual commissioning and the actual running system [8].

#### IV. DISCUSSION

The simulations which were performed showed the desired outcome. The simulation showed that the virtual device executed the operations and device movements in accordance with the PLC program. By obtaining this result, the PLC code is successfully validated. In addition, the simulations also provided a method to easily develop, improve, and troubleshoot the system PLC code.

With the physical commissioning, it was found that the predictions made through virtual commissioning, were successfully obtained. The only difference between the virtual simulation and physical commissioning was the speed at which it executed. This however can be rectified by changing the execution speed of the simulation and make the simulation more realistic; this will result in more accurate predictions and validation of the system.

#### V. CONCLUSION

The results obtained proved that virtual commissioning can be utilized to expedite the planning, verifying and optimizing of a system without the risk of damaging the real system equipment. It was possible to confirm the movements and operations of a device as well as validating the system PLC code through virtual commissioning in real-time. This shows that by using DELMIA to establish virtual commissioning, proper initial planning can be performed, apparent design faults can be resolved early in design stages, analysis can be done to validate changes to a system, and determine if it is actually profitable to build a system—thus build it right the first time or not at all.

To draw a conclusion, DELMIA provides a great tool for overall system verification and will ensure that manufacturing companies can validate their systems earlier, ramp-up production quicker and compete for global markets.

However, future considerations regarding virtual commissioning may include suggestive or predictive action from the virtual environment where the environment can propose solutions overlooked by a system analyst. Furthermore, features like automatically generating system code based on virtually validated systems will also be a welcome addition to virtual commissioning.

#### ACKNOWLEDGMENTS

The authors would like to thank the Central University of Technology Free State and AMTS under the project 07-11-P funding agreement; for the opportunity to conduct research on this topic.

#### REFERENCES

- [1] H. Krause, "Virtual commissioning of a large LNG plant with the DCS 800XA by ABB," in 6th EUROSIM Congress on Modelling and Simulation, Ljubljana, Slovénie, 2007.
- [2] N. Papakostas, G. Michalos, S. Makris, D. Zouzias, and G. Chryssolouris, "Industrial applications with cooperating robots for the flexible assembly," *International Journal of Computer Integrated Manufacturing*, vol. 24, pp. 650-660, 2011.
- [3] O. Salamon and A. Heidari, "Virtual commissioning of an existing manufacturing cell at Volvo Car Corporation using DELMIA V6," *Master Thesis, Department of Signals and Systems, Chalmers University of Technology, Sweden,* 2012.
- [4] H. Carlsson, B. Svensson, F. Danielsson, and B. Lennartson, "Methods for reliable simulation-based PLC code verification," *Industrial Informatics, IEEE Transactions on*, vol. 8, pp. 267-278, 2012.
- [5] Z. Liu, C. Diedrich, and N. Suchold, *Virtual Commissioning of Automated Systems*: INTECH Open Access Publisher, 2012.
- [6] A. Hollander and S. Sappei, "Virtual preparation of Tetra Pak filling machine," *Master Thesis*,

Department of Signals and Systems, Chalmers University of Technology, Sweden, 2011.

- [7] Kepware Technologies OPC Servers / Communications for Automation, Available Online: http://www.kepware.com, last accessed in August 2013.
- [8] J. Niemann, "Development of a Reconfigurable Assembly System with Enhanced Control Capabilities and Virual Commissioning," *Master Dissertation, Faculty of Engineering and Information Technology, Central University of Technology, Free State*, 2013.
- [9] 3DS Dassault Systemes, Available Online: http://www.3ds.com/products/delmia, last accessed in August 2013.
- [10] C. M. Park, S. M. Bajimaya, S. C. Park, G. N. Wang, J. G. Kwak, K. H. Han, et al., "Development of virtual simulator for visual validation of PLC program," in Computational Intelligence for Modelling, Control and Automation, 2006 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce, International Conference on, 2006, pp. 32-32.
- [11] S. C. Park, "A methodology for creating a virtual model for a flexible manufacturing system," *Computers in Industry*, vol. 56, pp. 734-746, 2005.
- [12] CATIA Online Documentation, Available Online: http://catiadoc.free.fr/online/CATIAfr\_C2/prtugCATI Afrs.htm, last accessed in August 2013.
- [13] D. Cachapa, A. Colombo, M. Feike, and A. Bepperling, "An approach for integrating real and virtual production automation devices applying the service-oriented architecture paradigm," in *Emerging Technologies and Factory Automation, 2007. ETFA. IEEE Conference on,* 2007, pp. 309-314.
- [14] H. Bley and C. Franke, "Integration of product design and assembly planning in the digital factory," *CIRP Annals-Manufacturing Technology*, vol. 53, pp. 25-30, 2004.
- [15] X. Kong, B. Ahmad, R. Harrison, Y. Park, and L. J. Lee, "Direct deployment of component-based automation systems," in *Emerging Technologies & Factory Automation (ETFA), 2012 IEEE 17th Conference on,* 2012, pp. 1-4.
- [16] C. G. Lee and S. C. Park, "Survey on the virtual commissioning of manufacturing systems," *Journal of Computational Design and Engineering*, vol. 1, pp. 213-222, 2014.
- [17] P. Hoffmann, R. Schumann, T. M. Maksoud, and G. C. Premier, "Virtual Commissioning Of Manufacturing Systems A Review And New Approaches For Simplification," in *ECMS*, 2010, pp. 175-181.

## **INDEX OF AUTHORS**

SURNAME	INITIAL	PAGE
Adewumi	A	207
Ajoodha	R	66
Akram	Μ	78
Azom	V	207
Badenhorst	J	118
Barnard	E	136
Becker	U	239
Behre	L	239
Benade	R	184
Berman	R	184
Bian	G	72
Booysen	т	195
Bosch	S	142
Bright	G	7, 37, 225
Burke	Μ	94
Cawood	G	83
Chikamai	К	30
Conning	Μ	19
Davel	Μ	118, 154, 166
de Villiers	J	24, 89
de Villiers	Н	130
de Wet	F	112, 136, 178
Dickens	J	148
Diegel	0	225
Donayo	J	7
du Toit	J	19
Dunbar	L	12
Els	S	100
Ferrein	A	172
Giwa	0	166
Gorlach	I	1, 83, 160
Govender	A	178
Green	J	60
Green	W	54
Grobler	Н	54
Haider	I	78
Harvey	Ν	245
Haskins	В	233
Helberg	A	142
Hofmeyr	D	201
lllidge	A	37
Jermy	R	24, 89
Kellerman	A	1
Khalid	H	78
Klein	R	66
Kleynhans	N	136
Kruger	G	72
Kuriakose	R	48

Limpert	Ν	172
Lipowicz	Μ	72
Lombard	н	100
Louw	J	252
Mangera	R	42
Marivate	V	106
Marnewick	Α	60
Mercorelli	Р	239
Modipa	Т	154
Morrison	G	42
Mzamo	L	142
Naidoo	Ν	7
Nel	Α	245
Nicolls	F	12, 24
Niemann	J	258
Niesler	Т	130
Nouhou	В	178
Pahl	С	213, 219
Pavlidis	Ν	201
Phillips	R	189
Poole	S	189
Pretorius	J	60
Reiner	F	195
Roberts	В	160
Rosman	В	66, 184
Sabatta	D	94
Sahraeian	R	112
Schiffer	S	172
Smith	S	19
Steele	Α	19, 123
Stopforth	R	7, 225
Stroud	Т	100
Swart	н	148
Tapamo	J	30, 207
Treurnicht	J	123
Uddin	V	78
Van Compernolle	D	112
van der Riet	D	225
van Niekerk	Т	160, 239
van Zijl	L	130
Varadarajan	К	219
Vermaak	н	48, 258
Viriri	S	30
Voigt	А	42
Walker	А	37
Werner	Ν	239
Z´arate	J	213
Zelasco	Z	7

Approved SITRAIN Partner

ry Automation

SIEMENS



# Siemens Accredited Training

The Advanced Mechatronic Technology Centre (AMTC) at Nelson Mandela Metropolitan University in collaboration with Siemens South Africa offers specialized Siemens accredited training, which include Service & Maintenance, Programming, Scada and Drives etc. The team at AMTC also do skills analysis / audits as well as structure tailor-made packages to meet your companies' needs.

For more info, please visit our website at <u>www.nmmu.ac.za/siemens</u> for a list and full description of the courses offered at the AMTC training Centre or contact us on 041 504 9912 for course bookings.





# Volkswagen South Africa - DAAD International Chair in Automotive Engineering

Leading new developments in energy systems, materials & manufacturing technology



# What can IEEE do for you on the job?

As IEEE members, professionals gain access to the deas, contacts and skill-enhancing resources that can accelerate development, as well as careers.

Discover more of what IEEE membership offers you. Visit our international online benefits tool. WWW.ieee.org/benefits







## Motion Control with SIMATIC S7-1500





For direct assistance contact: Siemens PE: Morne/Sakkie - 041 398 5000



# ASPIRE

# **Postgraduate studies**

Future dreams and ambitions are inherent in every one of us, but our future dreams are inspired by an initial individual aspiration, a starting point where you look at yourself and say 'yes I can, and I will achieve'.

With more than 320 postgraduate programmes across a vast selection of disciplines and over 200 short courses available, NMMU offers what you need to take a step closer to realising your ambitions.

#### Contact us today

- Andrew Kock
   Tel: 041 504 2162
- E-mail: postgrad@nmmu.ac.za



# Unbeatable Control, Precision, and Flexibility



Lowering cost, increasing productivity, and shortening design times are just some of the challenges industrial engineers face. The graphical system design approach combines productive software and reconfigurable I/O (RIO) hardware to help you meet these challenges. This off-the-shelf platform, customizable to solve any control and monitoring application, integrates motion, vision, and I/O with a single software development environment to build complex industrial systems faster.

>> Accelerate your productivity at ni.com/industrial-control-platform

South Africa: 0800 203199

©3014 National Instruments. All rights reserved. La W18W, National Instruments, NJ, and el.com are trademarks of National Instruments. Other productions of company names listed are trademarks or trademarks of their respective companies. I 1800 NI LabVIEW system design software offers ultimate flexibility through FPGA programming, simplifies code reuse, and helps you program the way you think-graphically.



