UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Department of Computer Science and Software Engineering

2019 Departmental Postgraduate Conference

September 5-6, 2019
Lecture Theatre 031, Jack Erskine Building

Principal Sponsor

verizon connect

Thursday 5 September (Jack Erskine room 031)

Session 1 *Chair: Richard Green* **Welcome, Keynote, CS Honours** (Judges: Andreas, Kourosh, Mukundan)	
9.10	**Welcome: Richard Green** **Keynote Speaker: Ben Reid** (AI Forum)
9:40	**Lucy Turner:** Maintaining Programming Skills of New Zealand Teachers
10:00	**Matt Smit:** Code style analysis for plagiarism detection
10.20	***MORNING TEA upstairs in the level 2 tea room***
Session 2 *Chair: Kourosh* **Masters** (Judges: Andreas, Kourosh, Mukundan)	
10:50	**Adam Tupper**: Harnessing Deep Neuroevolution to Create More Intelligent Game-Playing Agents
11:10	**Negar Mohammadhassan**: Towards Automatic Quality Assessment of Comments in Active Video Watching
11:30	**Timothy Irving**: Keypoint Detection of Tree Branch Using CNNs
11:50	**Jordan Gavan**: Finding the Characteristics of an Object of Interest Using 3D Modelling and Deep Learning
12:10	**Nicolas Zane Robinson-O'Brien:** Application of an algebraic framework to show correctness of Boruvka's minimum spanning tree algorithm.
12:30	***LUNCH upstairs in the level 2 tea room***
Session 3 *Chair: Kourosh* **PhD** (Judges: Tanja, Fabian, Mukundan)	
1:30	**Matthew Edwards:** Statistical Distribution of Checkerboard Corner Detection Error
1:50	**Caitlin Duncan**: An Intervention Study on the Development of Computational Thinking Skills in Primary School Students
2:10	**Camila Costa Silva**: Reusing software engineering knowledge from developer communication
2:30	**Sam Schofield**: Autonomous tree-pruning UAVs
2.50	**Sam Banks:** Moiré pattern interference used for pose estimation to aid home consumer grade control
3:10	***AFTERNOON REFRESHMENTS upstairs in the level 2 tea room***
Session 4 *Chair: Andreas* **PhD**	
3:40	**Sarmad Soomro**: A failed attempt to provide eyes-free interaction for in-vehicle touchscreen interaction
4:00	**Dibash Basukala**: Algorithms for Segmenting Substantia Nigra and Red Nucleus Regions in Quantitative Susceptibility Mapping Images
4:20	***Break for the day***

Friday 6 September (Jack Erskine room 031)

Session 5 ***Chair: Andy*** **PhD**	
9:10	**Keynote Speaker: Nathan Robinson (Verizon Connect)**
9:40	**Dilli Sharma**: Flexible Random Virtual IP Multiplexing for Moving Target Defense in Software-Defined Networks
10:00	**Ja'afaru Musa**: Using AVW to Support on the Job Training of Soft Skills for Software Industries
10:20	**Haipeng Li**: Micro-calcification Detection and Breast Density Classification Using Multifractal Analysis of Digital Mammograms
10:40	***MORNING REFRESHMENTS upstairs in the level 2 tea room***
Session 6 ***Chair: Tanja*** **PhD**	
11:10	**Rosalyn Rough:** Quantitative Bloodstain Pattern Analysis
11:30	**Bhagya Munasinghe**: How does CS Unplugged help teachers build confidence for teaching Computational Thinking (CT) topics?
11:50	**Faiza Tahir**: Extending Worked Examples in SQL-Tutor: from static to adaptive
12:10	**Tim Rensen**: Horticulture Related Deep Learning Using Synthetic Data
12:30	***LUNCH upstairs in the level 2 tea room***
Session 7 ***Chair: Richard*** **PhD**	
1:30	**Michelle Mashal**: Introduction to Deep Learning
1:50	**Bilal Ishfaq**: Network Level Defenses for Internet of Things (IoT)
2:10	**Amelia Samandari**: MAC Protocols for UAV Formations
2:30	**Dan Barry**: UAV Search Behaviour for Wireless Transmitter Detection
2.50	***AFTERNOON REFRESHMENTS*** **upstairs in the level 2 tea room**
Session 8 ***Chair: Fabian*** **PhD**	
3:20	**Oliver Batchelor**: Verification Based Annotation for Visual Recognition
3:40	**Kashif Amanullah**: Teaching Programming to School Children Using Elementary Patterns
4:00	**Josh McCulloch:** Conductor Reconstruction for Dynamic Line Rating using Vehicle-Mounted LiDAR
4:20pm	**Social gathering/Awards Ceremony in the Staff Club** Location: www.staffclub.canterbury.ac.nz/contact.shtml

Abstracts

Keynote Speaker

Ben Reid – Executive Director, AI Forum of NZ
Ben heads up the AI Forum of New Zealand - an independent organisation bringing together business, academia and the government connecting, promoting and advancing the AI ecosystem to help ensure a prosperous future New Zealand. Ben has a 24-year career background in software development and strategic technology consulting, working with a wide range of organizations spanning multinational software vendors, start-ups and government. Ben is active in promoting innovation and collaboration across the New Zealand technology industry. He was former Chairman of Canterbury Tech, the regional tech industry body and the establishment Chairman of Canterbury Angel Investors Inc., the regional early stage investment network. Ben is a member of the Institute of Directors in New Zealand.

CS Honours

Lucy Turner
Maintaining Programming Skills of New Zealand Teachers
Digital technology teachers in New Zealand have needed to up-skill themselves to teach the new DT curriculum, which includes computer science and computational thinking. Many of these teachers have rudimentary programming skills and go through periods of time where they are not using their programming skills. Based on related research it is likely that the programming skills of these teachers degrade when they are not being used (e.g. over the summer holidays) and that practising these skills regularly is likely to counteract this degradation. These findings support the need for codeWOF - a website developed by the UC Computer Science Education Research Group for teachers to regularly practise small programming exercises. I am also investigating whether the addition of Parson's problems or debugging exercises on top of the standard codeWOF exercises (where a program or function is written) impact programming skill maintenance.

Matt Smit
Code style analysis for plagiarism detection
All universities suffer from plagiarism, where students copy other students' work. In the field of computer science, software plagiarism can be a problem where students copy code from other students while solving assignments. Most work in this area has focused on similarity
detection, detecting how similar two works are, while taking into account different techniques plagiarists use to hide their attempts. Another technique is anomaly detection, finding features that occur in plagiarised work but not original work. We are researching code style analysis, forming a hypothesis that by looking into a students code submissions, we can detect anomalies that only occur when the student plagiarises.

Masters

Adam Tupper
Harnessing Deep Neuroevolution to Create More Intelligent Game-Playing Agents
In recent years, neural networks have been used to develop intelligent agents that have exceeded human performance in a variety of challenging tasks. To develop these agents, researchers have long used games as a testbed to evaluate and compare performance. Games are useful because they consist of a well-defined task, often require sophisticated strategies to master, and have validity in the sense that they are designed to challenge humans. Though gradient-based deep reinforcement learning methods have captured most of the attention recently, neuroevolution has also been shown to be a competitive method for training agents to master games and learn complex tasks. Neuroevolution refers to the design and training of neural networks using genetic algorithms and concepts based on the behaviour of natural evolution. In this talk, I will provide an overview of the field of neuroevolution for reinforcement learning, and discuss open questions and potential research avenues in the area.

Negar Mohammadhassan
Towards Automatic Quality Assessment of Comments in Active Video Watching
Learning by watching videos has been a conventional method of acquiring skills recently. AVW-Space is an online video-based learning platform for transferable skills which supports engagement via note-taking and personalised nudges. Measuring the quality of comments made by students in the note-taking environment can provide useful insights for analysing engagement. This study focuses on how to assess the quality of these comments. We first propose a scheme for assessing the quality of comments. Then, we evaluate this scheme by computing the inter-coder agreement and feature-based clustering. To automate the assessment of comments, we evaluate various machine learning approaches such as ordinal classification and cost-sensitive learning. The selected cost-sensitive classifier shows that the quality of comments can be assessed with relatively high F1-scores and low Mean Squared Errors. This study contributes to automating the assessment of short written reflections and the development of personalised educational support for engagement in video-based learning.

Timothy Irving
Keypoint Detection of Tree Branch Using CNNs
A project aims to prune pine trees in a commercial pine forests using a quadcopter UAV. A key part of this task is the detection of the cut point on a branch in a complex environment and in 3D space. Methods using CNNs on 2D images were developed. With this approach, a key point pixel is localised for the cut point and this is transformed into 3D space using RGBD cameras. It was seen in initial attempts that distraction from surrounding branches was a significant issue affecting results. The most successful method developed was a two-stage CNN architecture with a number of custom methods to mitigate distraction and improve keypoint accuracy. The CNN at the first stage was primarily focused on distraction and branch isolation, while the CNN at the second stage focused on accurate key point detection. A success rate of around 95% was achieved on a varied dataset.

Jordan Gavan

Finding the Characteristics of an Object of Interest Using 3D Modelling and Deep Learning

Using a depth-sensing camera to store depth and colour images from a variety of angles surrounding an object of interest, an object can be recreated in 3D space, using simultaneous localisation and mapping software, which includes background scene of the object from the images taken. The object can then be segmented from the 3D scene (background) using deep learning instance segmentation of the original images of the object to provide an object mask. The mask can then be projected onto the complete 3D scene to find the points which correspond to the object. The points are saved which results in a 3D object point cloud with the background removed. The 3D object can then be processed to find information such as dimension, volume or pose. This can be useful to a farmer as information such as livestock weight and condition of a fence can be found with almost zero human time involved.

Nicolas Robinson-O'Brien

Application of an algebraic framework to show correctness of Boruvka's minimum spanning tree algorithm.

Prior work has described an algebraic framework for proving the correctness of Prim's and Kruskal's minimum spanning tree algorithms. We intend to prove the correctness of another minimum spanning tree algorithm, Boruvka's, using the same framework. To this end, our results will be formally verified using the automated deduction capabilities of the Isabelle proof assistant. This work will further demonstrate the suitability of the framework as a sound abstraction for reasoning about weighted graph algorithms.

Sam Banks

Moiré pattern interference used for pose estimation to aid home consumer grade control.

Moiré patterns are the phenomena that occur as a result of interference between two similar periodic patterns. This interreference causes third interference pattern to appear. Small shifts to the original patterns can cause large shifts in the inference pattern or cause it to shift in unexpected ways. If the patterns are set at a distance apart this superimposed pattern becomes sensitive to changes in viewing angle and viewing distance. The goal is to use this effect to attempt to accurately estimate the elevation of a camera and ideally other pose information. Previous work has gone towards using moiré patterns formed with lenticular lenses to perform pose estimation for short ranges. However, positioning needs to be done accurately at large distances such as 15–30 metres and is too expensive to be suitable for grade control environments for the home consumer. Current work has gone toward developing theory to predict how the moiré pattern generated from two patterns at a distance changes with viewing angle and distance to test its practical feasibility. Preliminary tests from varying camera lateral translation appeared to be accurate for the close-range testing. Results from varying camera distance showed promise, however weren’t nearly as accurate as anticipated. Future work is predicted to be able to improve these early test results.

PhD

Matthew Edwards
Statistical Distribution of Checkerboard Corner Detection Error
The checkerboard pattern is one of the basic tools of computer vision. Since the late 90s, basically every computer vision researcher has used one for camera intrinsic calibration, and they're widely used for other purposes as well. However, little attention has been paid to their statistical properties. This presentation will go into great detail about the distribution of checkerboard corner detection error.

Caitlin Duncan
An Intervention Study on the Development of Computational Thinking Skills in Primary School Students
Computer Science (CS) and programming have been introduced to school curricula around the world, with the expectation that these will develop students' Computational Thinking (CT) skills. I have been investigating whether students' are learning and developing CT skills, through CS and programming curricula. Prior work conducted with Canterbury schools gave promising results, with the majority of participants reporting CS and programming lessons had positive impacts on students' learning and problem-solving skills. To follow up this study, an intervention study was conducted with three primary schools in 2017. Students took part in CS and programming lessons over 10 weeks of the school term. Their pre and post-test scores on CT tests were compared with control groups of students who didn't receive these lessons, and their teachers were interviewed. In this talk I present the results of this study, and the overall results of my research into CS and CT in schools.

Camila Costa Silva
Reusing software engineering knowledge from developer communication
Software development is a knowledge-intensive activity since it requires different types of knowledge, for example, knowledge about software development processes, practices and techniques and about the domain of an application. On the other hand, throughout the development of software, developers share knowledge via informal communication channels (e.g., instant messaging tools, e-mails, or online forums). Considering that this informal knowledge may be potentially relevant for other developers and given that this knowledge is not necessarily captured and formally documented for reuse, we investigate whether developer communication is a suitable source of reusable software development knowledge. Therefore, we aim to identify software engineering-related discussions from developer communications and evaluate their relevance for reuse in software development activities.

Sam Schofield
Autonomous tree-pruning UAVs
Forestry is the third-largest export industry in New Zealand. By pruning our commercial forests, we can produce more appearance grade timber, raising the value of each tree. Increasingly, the industry is struggling to find sufficient workers to carry out the physically demanding job of pruning trees. We aim to automate this task using swarms of UAVs. We have succeeded in pruning branches using a UAV in the lab, and we are working on moving into more realistic environments.

Sam Banks

Moiré pattern interference used for pose estimation to aid home consumer grade control

Moiré patterns are the phenomena that occur as a result of interference between two similar periodic patterns. This interreference causes third interference pattern to appear. Small shifts to the original patterns can cause large shifts in the inference pattern or cause it to shift in unexpected ways. If the patterns are set at a distance apart this superimposed pattern becomes sensitive to changes in viewing angle and viewing distance. The goal is to use this effect to attempt to accurately estimate the elevation of a camera and ideally other pose information. Previous work has gone towards using moiré patterns formed with lenticular lenses to perform pose estimation for short ranges. However, positioning needs to be done accurately at large distances such as 15–30 metres and is too expensive to be suitable for grade control environments for the home consumer. Current work has gone toward developing theory to predict how the moiré pattern generated from two patterns at a distance changes with viewing angle and distance to test its practical feasibility. Preliminary tests from varying camera lateral translation appeared to be accurate for the close-range testing. Results from varying camera distance showed promise, however weren't nearly as accurate as anticipated. Future work is predicted to be able to improve these early test results.

Sarmad Soomro

A failed attempt to provide eyes-free interaction for in-vehicle touchscreen interaction

Touchscreens are now commonly used to provide access to a wide range of vehicle functions. However, touchscreens require more visual attentional than physical control due to their lack of tangible sensation, creating important safety considerations. Prior research has proposed several feedback techniques to provide tactile sensation on a touchscreen, demonstrating reduced attentional demand while driving. In this study, we examine the use of physical augmentations (stencils) to provide eyes-free interaction while driving. We conducted an experiment in which two stencil designs were evaluated in comparison to a normal touchscreen during simulated driving. Results showed stencil overlays failed to provide eyes-free interaction while driving. We learned that the stencils we used in the experiment were not subtle to provide eyes-free interaction while diving; subjects relayed on visual feedback for target selection.

Dibash Basukala

Algorithms for Segmenting Substantia Nigra and Red Nucleus Regions in Quantitative Susceptibility Mapping Images

Substantia nigra (SN) and red nucleus (RN) located in midbrain are integral in the study of Parkinson's disease (PD). The automatic segmentation of SN and RN in high-resolution quantitative susceptibility mapping (QSM) images can aid in PD characterization and progression. However, only a few methods have been proposed to segment them, owing to the recent development of high quality imaging. Therefore, we describe a novel method for the segmentation of SN and RN in QSM images using contrast enhancement, level set method, wavelet transform and watershed transform. The segmentation performance is evaluated in 20 subjects consisting of both healthy and PD patients. The results of the proposed segmentation method were closer to the manual segmentation performed by the radiologist than the level set methods. The Dice coefficient of the left SN and right SN were

0.77 ± 0.09 and 0.78 ± 0.07 respectively while the Dice for the left RN and right RN were 0.80 ± 0.08 and 0.77 ± 0.08 respectively.

Keynote Speaker

Nathan Robinson – Technical Lead, Routing Algorithms at Verizon Connect
Conditional Contraction Hierarchies - Exact and Efficient Routing
Abstract:
The road network routing algorithm Contraction Hierarchies (CH) is able to answer queries up to six orders of magnitude faster than Dijkstra's algorithm. It achieves this by pre-processing a road network for a vehicle with specific dimensions and routing options (such as if toll roads are allowed). A major drawback of this approach is that we are unable to use this generated data to route optimally for vehicles with other dimensions. We propose an extension to CH that can route optimally for vehicles of any dimension and a range of routing options. We show that this algorithm performs comparably to the original CH algorithm on real-world continent sized road networks.
Bio:
I lead the routing algorithms squad at Verizon Connect and have been here for the past 4 and half years. Before this, I was a visiting researcher at NICTA and the Australian National University in Canberra, Australia where I worked on planning and propositional satisfiability. Before that, I was a post-doctoral fellow and lecturer at the University of Toronto. I did my PhD in automated planning at NICTA and Griffith University in Brisbane, Australia. I am interested in routing, planning, and optimisation, but logic (especially modal logic) and logical encodings of decision problems are my real passion.

Dilli Sharma
Flexible Random Virtual IP Multiplexing for Moving Target Defense in Software-Defined Networks
A moving target defense (MTD) is a proactive dynamic defense strategy that continuously changes the attack surface. Network address shuffling is one of the MTD techniques that invalidates the address information collected by the attacker. We proposed a shuffling-based MTD technique for software-defined networking (SDN) called Flexible Random Virtual IP Multiplexing, namely FRVM, which aims to defend against network reconnaissance and scanning attacks. FRVM enables a host machine to have multiple, random, time-varying virtual IP addresses, which are multiplexed to a real IP address of the host. Multiplexing or de-multiplexing event dynamically remaps all the virtual IP addresses of the hosts. Therefore, at the end of a multiplexing event, FRVM invalidates any knowledge gained by the attacker through the reconnaissance. We evaluated the performance of the FRVM in terms of the attacker success probability under scanning attacks and operational overhead.

Ja'afaru Musa
Using AVW to Support on the Job Training of Soft Skills for Software Industries
Numerous research projects have focused on investigating the approach for improving formal education (e.g., in schools and universities), while little attention has been concentrated towards providing soft skills development in the industry (such as software development industry). Soft skills (e.g., the ability to collaborate, communicate and

negotiate) have proved to be an important complement to technical skills (e.g., the ability to analysis, design and code or develop) in the current software development environment. Nevertheless, instructing soft skill in a customary educational setting is costly and time consuming for small and medium software companies. In this study we aim to investigate the efficiency of using Active Video Watching (AVW) as a new approach towards incorporating the training of soft skills into industrial education as an alternative to the traditional resource-intensive and time-consuming approach.

Haipeng Li

Micro-calcification Detection and Breast Density Classification Using Multifractal Analysis of Digital Mammograms

Mammography is a commonly used method for breast cancer screening. Computer-aided Detection (CADe) methods are developed to assist radiologists with the detection and segmentation of regions of interest and quantification of disease characteristics. This work focuses on the development of algorithms for micro-calcification detection and breast density classification in digital mammograms, using multifractal analysis combined with machine learning and image processing methods. Experimental analysis shows that texture features based on multifractal descriptors are enhanced significantly using the proposed method and a good detection result is observed with a SVM classifier and a CNN-based false positive reducing method. For breast density classification, multifractal spectrum is used to characterise local fibro-glandular tissue textures, combining with histogram analysis and Chi-square classifier, which shows a better classification result.

Rosalyn Rough

Quantitative Bloodstain Pattern Analysis

This research project is a multi-discipline project which intends to apply computer vision and deep learning techniques to the forensic discipline of bloodstain pattern analysis. This will build on prior research which innovated computer vision based methods to identify and annotate bloodstains in a digital image. Metrics of the individual stains and the pattern (whole collection of stains) are also to be identified. A ground truth dataset of known patterns will be generated and the prior research will be extended to assist with establishing discriminatory quantitative features in each pattern type. Furthermore, it is proposed that deep learning neural networks ResNet and PointNet will be adapted to automatically classify the bloodstain patterns, where the images from the dataset of known patterns, and the data mined from them will be used to train such adapted networks.

Bhagya Munasinghe

How does CS Unplugged help teachers build confidence for teaching Computational Thinking (CT) topics?

Countries worldwide have incorporated CT in school curriculums, introducing it to students as young as 5 years old, with intentions of changing them from users and consumers of digital technologies to skilled innovative creators of digital solutions. This introduction requires teachers to be contend with teaching of new Computer Science (CS) content. Finding effective ways to teach CT and deliver CS has been an active discussion and the smaller number of teachers having computing background has been an issue. Computing teachers need pedagogical content knowledge (PCK), which includes awareness of common misconceptions, methods for diagnosing them, and interventions to help students develop

more robust conceptions. Teachers' feedback has consistently suggested unplugged as a successful method for teaching CT and introducing programming concepts to young students. Having that in mind, we intend to research how CS Unplugged can help teachers build confidence for teaching CT.

Faiza Tahir

Extending Worked Examples in SQL-Tutor: from static to adaptive

Advanced learning technologies particularly Intelligent Tutoring Systems (ITSs) are playing a vital role in today's learning environment. The goal of these systems is to provide individual support to learners according to their needs and abilities. Previous work done on SQL-Tutor, an ITS that teaches database querying, provided problem-solving opportunities to student adaptively. My PhD project aims to develop an adaptive strategy for providing examples during problem solving. The strategy would need to identify the right time and right type of example during problems solving. The study to answer these questions may include estimation of the learner's knowledge and their behaviour towards using examples, identification of affective states of learners and their motivation during problem solving and particularly while working with examples, strategies to improve motivation and affect states and strategies to improve example using behaviours and finally the effects of using different types of examples in combination of these strategies.

Tim Renson

Horticulture Related Deep Learning Using Synthetic Data

For the purpose of training a convolutional neural network a synthetic data set of "plant" images with corresponding node point heatmap annotations was generated. Visualisation Tool Kit (VTK) was used to generate triangle surface mesh to represent leaves and fruit, and spline tubes were used to represent the stem structure. When trained on a dataset of 1000 random angle image/anotation pairs of two randomly generated plants, a simple convolutional neural network was able to create a node heatmap on test data from a third randomly generated plant. The goal of this research is to train neural networks on large datasets of simulated data, then finetune the pretrained network on real images. An ideal network will function as a map between stereo images and a simple 3D representation of a plant for agricultural and horticultural robotic applications.

Michelle Mashal

Introduction to Deep Learning

Over the past decade making prediction using deep learning created outstanding results in operation efficiency in many industries. Unlike typical data analysis, deep learning prediction models comprises of large, diverse and complex data. Deep learning is a sub-field of machine learning that has revolutionized several fields such as image processing, speech recognition and natural language processing. Due to success in these areas deep learning has gained a lot of attention from the scientific community.

Even though deep learning has been used successfully in many industries, so far there has been hesitation in using deep learning prediction model in aviation due to several open problems that needs to be solved in order to make major breakthrough in aviation using deep learning. Initialising and training a deep learning model requires expert knowledge since there are many open problems still exists with deep learning. One of the main problems is that there isn't complete understanding of why prediction with deep learning

works so well. This research aims to use the Airways New Zealand's complex aviation systems data and find case studies in aviation fields to explain how to use deep learning prediction models to connect to the data from controlled airspace. Using the current knowledge of deep learning and known limitations the aim is to refine the prediction methods and develop a new prediction model ideal for controlled airspace future operation prediction. Keywords: Controlled Airspace Prediction, Deep learning

Bilal Ishfaq
Network Level Defenses for Internet of Things (IoT)
Internet of Things (IoT) has become a point of attraction to the industry and academia recently. IoT network consists of heterogeneous devices with limited computational and power resources, and not equipped to dynamically respond to abnormalities, eventually exposed to various cyber-attacks. To ensure the security of IoT networks, we propose an approach which aims at the design, implementation and testing of network-level defense techniques for IoT. We propose to implement a small-scale Software-Defined IoT (SD-IoT) testbed and collect security related data from the testbed. We use statistical inference techniques to measure the security related data including mean time to compromise an IoT node, mean time to detect an intrusion, mean time to reconfigure the SD-IoT, etc. We also design, implement and test the reactive and proactive defense mechanisms for IoT which includes the moving target defense (MTD) techniques.

Amelia Samandari
MAC Protocols for UAV Formations
UAV formations can be used to support diverse industries, with applications in entertainment, natural environment, civilian and military spheres. Using UAV formations is more effective than deployment of a single UAV, because multiple UAVs cooperating to complete a mission are able to perform more complex tasks with greater precision and efficiency.

We wish to provide Quality of Service for safety messages through a MAC protocol that appropriately addresses the issue of packet collisions for inter-UAV communication. UAV formations can then be more effective and successful in completing the tasks required, irrespective of the industry or application in question.

Dan Barry
UAV Search Behaviour for Wireless Transmitter Detection
With the popularization of unmanned aerial vehicles (UAVs) in both a commercial and hobbyist setting, UAVs have been used in increasingly varied fields due to reduction in cost and increased capability. Our work focuses on the problem domain of search and rescue (SAR), where we consider a scenario that an unknown number of victims are to be discovered that carry small unknown transmitters, in an unknown environment that varies in complexity from an open field to a dense urban area. We look to reduce both the time taken to find all victims and time taken to reach some confidence threshold about additional victims within the search area. Through the use of a challenging discretized simulation environment, the UAV agents are required to implement a behaviour that allows them to both explore and search simultaneously. Our approach is motivated by the information theoretic concept of empowerment, where agents consider large numbers of future action

sequences that maximize information that can acquired for their internal model of the environment.

Oliver Batchelor
Verification Based Annotation for Visual Recognition
Applying machine learning to new domains in visual recognition usually implies a data collection and annotation problem. I have developed and evaluated human-in-the-loop methods for image annotation based around verifying machine predictions, for the purpose of rapid prototyping and experimentation. I demonstrate the effectiveness of the method by annotating a variety of object detection data-sets, as well as applications to verification based counting of wildlife.

Kashif Amanullah
Teaching Programming to School Children Using Elementary Patterns
Teaching programming to school children is a challenging task. Block-based programming languages serve as a great tool for introduction to programming but not without problems. There is a dire need of a better teaching pedagogy for teaching programming, and elementary patterns present themselves as a structured approach. This research looks at the issues with block-based languages and suitability of elementary patterns as a possible solution. Results to date show a little use of important programming elements even after many years of use, and there is no clear sign of progression in skills with or without remixing.

Josh McCulloch
Conductor Reconstruction for Dynamic Line Rating using Vehicle-Mounted LiDAR
Dynamic Line Rating (DLR) is a process which electrical network operators can implement to improve efficiency by dynamically adjusting the load capacity as conditions allow. To implement DLR an accurate model of the conductors and their clearances is needed. Airborne LiDAR, while expensive, is the most common method of collecting line data as it is fast and is of high quality. State of the art methods for automatically reconstructing conductors first classify conductor points before fitting conductor models. This approach works well for high tension lines with significant separation between conductors but tends to perform poorly urban environments where conductors are packed tightly together and surrounded by clutter. The presented method attempts to overcome these challenges by performing an informed search for the conductors, anchored to the utility poles. Before the conductors are classified, their layout and sag are estimated, converting conductor segmentation into a linear problem, and a 3D to 2D projection is used to improve density and simplify clustering. The work also attempts to reduce the cost of conductor reconstruction by utilising lower-cost vehicle-mounted LiDAR. By avoiding point classification, higher precision can be achieved in scenarios where previous methods have suffered from significantly degraded performance.

Thanks to our Sponsors: