Department of
Computer Science and Software Engineering

2023 Postgraduate Conference
Thursday 7th and Friday 8th September 2023
JE031, Jack Erskine Building

Principle Sponsor:

verizon✓
connect

Also sponsored by:

Thursday 7 September

Session 1	*Chair: Tim Bell*
9:30 – 10:00	**Keynote presentation: Dr Nathan Robinson, Marat Sibaev – Verizon Connect** *The Routing Cost Model*
BSc Hons (Judges: Tim, Mukund, Fabian)	
10:00-10:15	William Valentine (BSc Hons) *GPU Accelerated Modelling and Real-time Rendering of Fluid Motion*
10:15 – 10.30	Gareth Harcombe (Bsc Hons) *Locating Athletics Tracks from the GPS Data of a Runner*

10:30 – 11:00	***Morning tea** (upstairs in the level 2 tea-room)*

Session 2	*Chair: Aluna Everitt*
Masters (Judges: Tim, Mukund, Fabian)	
11:00 – 11:15	Drew de Wet (ME) *Development of a path planning algorithm for a robotic arm to prune grape vines*
11:15 – 11:30	Jamie Houghton (MSc) *Integrating Remote Sensing Datasets for Global Scale Flood Risk Mapping*
11:30 – 11:45	Dylan White (ME) *A filtering approach to online monocular camera calibration*
11:45 – 12:00	Tim Chang (MSc) *Hierarchical Reinforcement Learning in the Game of Othello*
12:00-12:15	Yuanjie Shi (MSc) *Traffic Sign Detection and Recognition*
12.15-12.30	Marina Filipovic *Predicting project success based on testing metrics in a year-long project*

12:30 – 1:30pm	***LUNCH** (upstairs in the level 2 tea-room)*

Session 3	*Chair: Mukundan*
PhD (Judges: Kourosh, Andrew, Etienne)	
1:30 – 1:45	Amelia Samandari *TDMA Slot Allocation for UAV Formations*
1.45 - 2.00	Andrew Davidson *Machine Learning Classification of Histologic Information of RNAscope Stained Breast Cancer Tissues Using Whole Slide Images*
2:00 – 2:15	Harry Dobbs *Smart-Tree: Robust 3D Tree Skeletonization from Point Clouds*
2:15 – 2:30	Casey Peat *Vineyard 3D Reconstruction - A Complete Pipeline*

2:30 – 3:00	***Afternoon tea** (upstairs in the level 2 tea-room)*

Session 4	*Chair: Fabian Gilson*
3:00 – 3:15	Gus Ellerm *LivePublication: The science workflow creates and updates the publication*
3.15 – 3.30	Aaron Smith *Neural Radiance Field Spectral Computerised Tomography Reconstruction*
3:30 – 3:45	Zhouyu Qu *Wireless path planning algorithms on a realistic drone road system*

Friday 8 September

Session 5	Chair: Andrew Bainbridge-Smith
10:15 – 10:30	Ehsan Bojnordi *Adding Interactive Mode to Active Video Watching*
10:30 – 10:45	Sreedevi Iyer *Empathy measure for software professionals*
10:45 - 11:00	Pasan Peiris *Learner Perceptions on Gamifying Active Video Watching Platforms*

11:00 – 11:30	***Morning tea*** *(upstairs in the level 2 tea-room)*

Session 6	Chair: Kourosh Neshatian
11:30 – 11:45	Sujan Warnakulasooriya *Novel method for optimising initial drone flocking*
11:45 – 12:00	Raul Lumapas *Question-Driven Design Process for XAI in Active Video Watching*
12:00 – 12:15	Ben McEwen *Active Few-shot Learning for Sparse Bioacoustic Feature Annotation*
12:15 – 12:30	Tim Rensen *Photogrammetry and machine learning for surveying scallops*

12:30 – 1.30pm	***LUNCH*** *(upstairs in the level 2 tea-room)*

Session 7	Chair: Trevor Nesbit
1:30 – 1:45	Sharon Temtsin *A Bona fide Turing Test*
1.45 - 2.00	Matthew Minish *Promoting Student Self-reflection in Software Engineering Project Courses through Digital Check-ins*
2:00 – 2:15	Henry Hickman *Beyond Question Shuffling- Randomisation in Programming Assessment*
2:15 – 2:30	Sam Schofield *End to end machine learning for robust robotics?*

2:30 – 3:00	***Afternoon tea*** *(upstairs in the level 2 tea-room)*

Session 8	Chair: James Atlas
3:00 – 3:15	Juliet Samandari *Post-quantum Authentication in MQTT*
3:15 – 3:30	Nick Lee *Visual Odometry and Pose Estimation Error*
3:30 – 3:45	Rosalyn Rough *Improving the Reliability and Ease of Bloodstain Pattern Analysis with Quantitative Techniques*

4:00	***Staff Club for prizegiving and celebration*** *Location:* **https://www.ucc.org.nz/contact-us**

Abstracts

BSc Honours

Gareth Harcombe

Title: Locating Athletics Tracks from the GPS Data of a Runner.

Abstract:
High Performance track runners are in need of tools to analyse their workouts, such as measuring lap times, identifying maximum acceleration, and predicting race performance. However, the data collected from GPS watches is not well suited for this task, as it is missing features such as where the runner is on an athletics track. This feature selection requires the location of the athletics track that the runner is running on, and must be done with the constraints of the limited GPS data that a standard GPS watch collects, GPS inaccuracies, and a small dataset. We propose three methods to find the coordinates of an athletics track from a runner's GPS data: Convolutional Neural Networks; traditional Computer Vision techniques; and numerical optimisation.

William Valentine

Title: GPU Accelerated Modelling and Real-time Rendering of Fluid Motion

Abstract:

Simulating fluids is an important task in engineering and scientific research contexts as well as in video games and animation. We propose a fluid rendering pipeline that uses OpenGL-4 shaders to employ the parallel processing capabilities of the GPU. The fluid's surface mesh is produced using tessellation shader stages where the input patches are assigned tessellation levels based on the fluid heightmap's curvature. The curvature is stored using a texture buffer object which allows access by shaders, thus allowing the tessellation calculations to be carried out in parallel. Use of this adaptive tessellation method increases both the simulation's framerate as well as its capacity to handle a greater number of primitives. Furthermore, it more optimally distributes the mesh's vertices to effectively increase the level of detail without using more primitives. Polygon culling using the geometry shader further optimises the number of primitives used to define the fluid surface. We also propose two GPU-based fluid surface flow visualisation methods. Texture buffer objects can be used to store a surface texture that is updated each frame. Alternatively, particle positions are updated each frame using the geometry shader and stored in a buffer object using transform feedback. These flow visualisation techniques are particularly effective for communicating the swirling motion of vortices.

Masters

Dylan White

A filtering approach to online monocular camera calibration

Abstract: Structure-from-motion, visual odometry, SLAM and many other computer vision technologies rely on the ability to interpret the physical geometry of a scene captured by a camera. To extract accurate geometric information from the captured scene, the camera must be calibrated to establish a robust estimation of the projection function from the scene to the sensor. This is typically a laborious task. If done in the factory, it adds to the cost of manufacturing. If done by the user, it requires expertise and specialized equipment that makes the system less generally useful. Furthermore, calibrations are not stable over time - they change with age, temperature, shock, and vibration. The traditional method of camera calibration, which is completed offline as a preprocessing step using a predefined calibration pattern, is susceptible to shifts in calibration over time which can degrade the accuracy of measurements.

Our research proposes a new method for online monocular camera calibration which utilizes data from an IMU and GNSS receiver rigidly attached to the camera to jointly estimate the camera pose, camera intrinsic parameters, and feature positions in real-time. This allows for continuous detection and correction of any shifts in the calibration. We formulate this as a state-space problem and use high rate IMU data to predict relative camera and feature motion, while GNSS observations and visual data tracking using sparse optical flow serves as the update. The inference is solved using an unscented Kalman filter implementation. Unscented Kalman filtering is a widely used approach in the field of visual-inertial sensor fusion for estimating the state of a dynamic system. It provides a means of incorporating information from multiple sources, such as visual and inertial sensors, to improve the accuracy and robustness of state and parameter estimation.

As Trimble continues to integrate vision technology into its solutions, our proposed approach for online camera calibration becomes increasingly relevant for ensuring long-term calibration stability for users.

Drew de Wet

Development of a path planning algorithm for a robotic arm to prune grape vines

Abstract: Sampling based path planning algorithms are widely used algorithms that find a feasible path for a high dimensional robot to move to a goal pose. They allow for path planning in environments with complex collision objects such as grape vines by finding a path to the goal between random collision free samples in the robot's configuration space. A simulation environment is used to test and develop hardware and algorithms. The UF850 robotic arm is selected and RRT* is found to be an algorithm that provides a good approach to be adjusted and implemented in a way that suits the application. The path planning algorithm and related software and mechanical systems are developed towards the goal of fully autonomous pruning being executed reliably and efficiently.

Jamie Houghton

Title: Integrating Remote Sensing Datasets for Global Scale Flood Risk Mapping

Abstract: The need to accurately predict flood risk at a global scale is urgent. Population exposure to flood hazard is increasing and flooding impacted 2 billion people globally over the last two decades according to the Centre for Research on the Epidemiology of Disasters (CRED). Flood risk maps typically assign a probability of flooding for each location on an annual basis. However, this estimate is limited by temporal coverage and cloud cover during flooding. We work in human terms to predict the annual number of days of flooding at a 1km grid at a global scale by developing a model to more accurately identify historic flooding using an integration of remote sensing datasets. Landsat imagery provides global coverage and Dartmouth Flood Observatory (DFO) provides a database of large flood events, both spanning the previous four decades. The Normalized Difference Water Index is used to distinguish water with change detection to denote flooding. We integrate precipitation, river discharge, soil moisture, and other layers to train a model to identify likely flooding when imagery is not available to detect floods through remote observation. Increasing data collection and aggregation at the global scale allows for the use of data intensive machine learning models to analyze these factors for the task of filling gaps in flood risk mapping.

Marina Filipovic

Title: Predicting project success based on testing metrics in a year-long project

Abstract: Testing plays an important role in securing the success of a software development project. Various techniques of automated testing have been developed, including automated acceptance testing which represents the customer's expectations in such a way that the requirements can be translated into automated tests.
In this research, we investigate whether we can quantify the effects various types of testing have on functional suitability, i.e. the software conformance to users' functional expectations. Additionally, we investigate which combination of testing (automated and manual) should be done to ensure the expected requirements of software are met, as well as whether the lack of testing in the early stages of the project causes a significant increase of effort spent fixing the project later on?

To answer those questions, we collected and analysed data from a year-long software engineering project course. The results suggest that the level of effort teams put into the testing that is the closest to the customer's experience (i.e. manual & automated testing at the user interface level) has a significant impact on functional suitability. We also found a statistically significant relationship between the total amount of effort spent on testing in one sprint and fixing effort in the following.

However, based on the collected data, automated acceptance testing does not seem to have a strong influence on product quality, even though its purpose is to translate user's expectations into tests. Still, the literature suggests that automated acceptance testing has a beneficial impact on functional suitability, therefore, we believe that these results can be explained with the students' lack of experience with more advanced automated testing techniques and recommend providing students with additional guidance and training opportunity in that regard.

Timothy Chang

Title: Hierarchical Reinforcement Learning in the Game of Othello

Abstract: My thesis is on Hierarchical Reinforcement Learning in the game of Othello. In my thesis, I have created a hierarchical reinforcement learning algorithm that plays Othello and comparing it's sample efficiency, information gain and transferability to normal reinforcement learning. I found that hierarchical reinforcement learning requires significantly less data, training time and experience to create an AI for Othello than using standard reinforcement learning algorithms. I have also created a new method of hierarchical reinforcement learning algorithm that can be used for other applications.

Yuanjie Shi

Title: Traffic Sign Detection and Recognition

Abstract: In this research project, a machine learning-based pipeline that is used to recognize and classify traffic signs would be built based on an existing innovation project owned by Verizon Connect. With this pipeline, data captured from truck-based cameras would be used as the input, and the type of traffic sign and its geo-information would be the output. This input video data will be sourced from the Berkeley Deep Drive video library, licensed by Verizon Connect. Deep learning techniques will be implemented in this project, and models with candidate architectures would be measured by accuracy rate, information density, and NetScore. We expect to use GTSRB, the German Traffic Sign Benchmarks (CC0 1.0 license, i.e. copyright free) for validation.

PhD

Aaron Smith

Title: Neural Radiance Field Spectral Computerised Tomography Reconstruction

Abstract: Computerised Tomography (CT) reconstruction is the act of generating a 3D volume from a series of 2D projections typically captured in a helical pattern of a CT scan, which is a special case of an image mapping problem. Most modern CT reconstruction algorithms are based on a model-based reconstruction method where the representation of the object in the scanner is repeatedly updated and tested to analyse the accuracy of the generated model. We propose the adapting and utilisation of a Neural Radiance Field (NeRF) network as an alternative model-based CT reconstruction approach.

Amelia Samandari

Title: TDMA Slot Allocation for UAV Formations

Abstract: This presentation looks at the case of mitigating collisions between Unmanned Aerial Vehicles (UAVs) in a formation through the use of safety beacons to relay information about UAVs that are at risk of collision due to their geographic proximity to one another. The UAVs send these safety beacons using time division multiple access (TDMA). With TDMA, UAVs can achieve collision-free transmission, thereby reducing the uncertainty of the UAVs in the formation receiving the necessary safety information. Through spatial reuse of TDMA slots, theoretical results demonstrate that our scheme can support large formations with a bounded transmission period.

Andrew Davidson

Title: Machine Learning Classification of Histologic Information of RNAscope Stained Breast Cancer Tissues Using Whole Slide Images

Abstract: The genetic characteristics of breast cancers determine the mechanism that the cancer uses to grow, and therefore, the treatment options that will be effective. RNAscope in situ hybridization allows for the staining of particular RNA sequences in tissue samples, and therefore the quantification of gene expression. However, accurately quantifying the level of RNAscope staining in tumour and non-tumour tissue is extremely time-consuming when done manually. My research aims to develop whole slide image analysis techniques and machine learning algorithms that will provide accurate quantification of gene expression. This will allow for easier and more objective appraisal of RNAscope stained breast cancer tissues.

Ben McEwen

Title: Active Few-shot Learning for Sparse Bioacoustic Feature Annotation

Abstract: The collection and annotation of bioacoustic data presents a number of challenges to researchers. Focus is often constrained to highly vocal species and limits the study of less vocal species. We present Listening Lab, a tool designed to aid in the analysis of sparse audio events within long-term field recordings. This tool combines a wavelet-based segmentation method with a few-shot active learning recommender system in a human-in-the-loop process to prioritise the annotation of low-certainty samples. This process combines the accuracy of human classification and the speed of computational tools to greatly reduce the empty space in sparse field recordings. This tool allows users to train custom audio classification models for any application with sparse features and this classification

pipeline can be easily exported for use in the field making real-time bioacoustic monitoring of less-vocal species a possibility.

Casey Peat

Title: Vineyard 3D Reconstruction - A Complete Pipeline

Abstract: For the purpose of autonomous vine pruning, we present a complete pipeline for building a 3D reconstruction of the vineyard scene. This includes our novel additions to NeRF to accurately extract the 3D geometry, our method Zero-NeRF to align 3D scenes with "Zero" Overlap, and our error detection and correction methods to make this pipeline a viable solution in practice.

Ehsan Bojnordi

Title: Adding Interactive Mode to Active Video Watching

Abstract: Presentation skills are crucial for tertiary students and graduates but are difficult to teach. We augmented active video watching (AVW) approach with the possibility of interactions among students, and conducted an experiment with AVW-Space, an online platform which supports video-based learning. Participants of the experiment watched and commented on videos first. In the second phase, the participants reviewed, rated, and responded to their peers' comments. We found that students who interacted with other students and responded to their comments increased their conceptual understanding of presentation skills.

Gus Ellerm

Title: LivePublication: The science workflow creates and updates the publication

Abstract: The uptake of computational methods to support research has led to some remarkable new tools and methods to improve outcomes. But one unintended consequence is that the scientific record ends up being fragmented and distributed amongst several distinct systems. The research we report aims to gather together **all** of the components of an experiment into a single container---including the publication itself. We describe the architecture of such a system that marries together distributed workflows (Globus) with research object containers (RO-Crate) and adds new methods to describe, update and `publish' the details of the workflow and its outcomes. Finally, we demonstrate the system with a natural language processing research use case.

Harry Dobbs

Title: Smart-Tree: Robust 3D Tree Skeletonization from Point Clouds

Abstract: This presentation introduces Smart-Tree, a supervised method for estimating tree branch skeletons' medial axes from point cloud data. We employ a sparse voxel convolutional neural network to extract radius and direction information for each input point, followed by a greedy algorithm for skeletonization.

Smart-Tree addresses challenges like complex tree structures, self-occlusions, intricate geometry, intersecting branches, and varying point cloud densities. We evaluate it using synthetic and real-world datasets, showcasing its robustness compared to existing methods.

Henry Hickman

Title: Beyond Question Shuffling- Randomisation in Programming Assessment

Abstract: Randomisation is a technique that can be used with programming assessments to discourage academic misconduct by making it unlikely for two colluding students to get the exact same questions. Previous research about randomisation has shown it to be an effective tool for addressing academic misconduct, but this work often focuses on randomisation broadly, with few considering specific techniques. Here we consider specific randomisation techniques and the contexts that they are best suited to. In addition, we investigate the effectiveness of randomisation techniques against emerging AI technologies. This is done with the goal of implementing these techniques in an automated assessment system to assess programming in New Zealand high schools.

Juliet Samandari

Title: Post-quantum Authentication in MQTT

Abstract: Message Queue Telemetry Transport (MQTT) is a common communication protocol used in the Internet of Things (IoT). MQTT is a simple, lightweight messaging protocol used to establish communication between multiple devices. However, the protocol does not provide authentication, and most proposals to incorporate it lose their lightweight feature and do not consider the future risk of quantum attacks. IoT devices are generally resource-constrained, and postquantum cryptography is often more computationally resource-intensive compared to current cryptographic standards, adding to the complexity of the transition. We will discuss the options and considerations for incorporating postquantum security into this protocol.

Matthew Minish

Title: Promoting Student Self-reflection in Software Engineering Project Courses through Digital Check-ins

Abstract: Software engineering education relies heavily on experiential learning (e.g. group project courses) for students to learn skills that cannot be taught effectively through lectures alone. Self-reflection is a key ingredient in learning from experience, and as such is a highly important practice in experiential learning.

In the 2022 iteration of the SENG302 group project course run at UC, we introduced an intervention aimed at promoting self-reflection of students during their everyday working process. This intervention took the form of a pre-stand-up meeting digital check-in, where students reflected on how they had approached their assigned tasks since the previous stand-up, including any lessons learned or obstacles encountered. In this presentation we report on our experience using this intervention in 2022, in addition to discussing the changes made for the continuation of the study in the 2023 instance of the course.

Nick Lee

Title: Visual Odometry and Pose Estimation Error

Abstract: Visual odometry estimates the robot's pose by analysing images captured by it and is integral to autonomous navigation. However, when devices operate outdoors, the large changes in brightness may cause localisation failure. This short-coming motivates us to find methods that ensure visual odometry can function in high dynamic range environments. To offset the negative impact of operating in high dynamic range environments, researchers proposed automatic exposure algorithms to select a

camera exposure time that maximises their selected quality metric. However, direct comparisons between the selected image metric and visual odometry performance is scarce. Many of the proposed algorithms go through multiple steps, making it difficult to determine whether the improved pose estimation accuracy is the result of optimising the metric.

In this submission, we generated sets of synthetic images in a simulation environment at different exposure levels. We estimate trajectories ten times for each data set and compare the resulting errors with the image quality metrics. Our tests showed that none of the selected image quality metrics had a strong, negative correlation with absolute translation error. The findings suggest that the current automatic exposure algorithms have yet to find an image quality metric that truly impacts visual odometry performance. Ongoing research should consider alternative metrics or metric combinations.

Pasan Peiris

Title : Learner Perceptions on Gamifying Active Video Watching Platforms

Abstract: Video-based learning (VBL) provides self-paced and flexible learning. However, VBL is often a passive learning method. Active video watching (AVW) has been proposed as an approach to increase learner engagement. We investigate the motivation and perception of learners towards gamification to further increase engagement in AVW. Results from a survey in New Zealand and the Philippines show a positive perception towards integrating gamification into AVW, with learners preferring a combination of game elements rather than individual elements. Our findings provide foundations for a gamification intervention in AVW.

Raul Lumapas

Title: Question-Driven Design Process for XAI in Active Video Watching

Abstract:
Designing explanations for Artificial Intelligence (AI) systems continues to be a challenge due to AI's black-box nature. Among the solutions developed to help in designing explanations in AI technologies is the Question-Driven Design Process for Explainable Artificial Intelligence (XAI) User Experience. In this paper, we report on our experiences using the question-driven design process for XAI in active video watching. We used Active Video Watching (AVW)-Space, an AVW platform developed at the University of Canterbury, as the context for AVW. We conducted a survey to elicit questions from users on the AI features of the system. In the question analysis process, three human raters categorized the user questions into the different XAI bank categories. Results show that most users tend to ask "how" and "why" questions about the AI-enabled features in the platform. The results of the question analysis will be used in mapping the determined question categories to potential XAI techniques. This can help in deciding the types of explanations to provide to users of AVW in future works on XAI in active video watching.

Rosalyn Rough

Title: Improving the Reliability and Ease of Bloodstain Pattern Analysis with Quantitative Techniques

Abstract: Bloodstain Pattern Analysis (BPA) is a discipline of forensic science which is often used at scenes of bloodshed to assist with the reconstruction of events. Seen as a pattern recognition discipline, BPA has received criticism of its subjective nature and lack of quantitative techniques. As a result, quantitative classification methods are being explored in this research project.

The development and validation of an automatic bloodstain pattern method will be discussed. This method utilises computer vision techniques to identify bloodstains on a plain background within a digital

image and generate bloodstain specific metrics relating to the individual stains as well as the overall pattern This method improves on prior work with regard to the segmentation of the individual bloodstains within an image and also returns more useful discipline-specific metrics. The application of the method to digital images of bloodstain patterns allows analysts to obtain much more quantitative data than is currently possible using manual techniques or other image-based programs currently utilised in the discipline. This method has been applied to a dataset of laboratory generated impact, cast-off and expired patterns, yielding a wealth of quantitative data.

Sam Schofield

Title: End to end machine learning for robust robotics?

Abstract: Robots are often more precise than humans at both sensing and actuation yet are typically much less robust at navigation and manipulation tasks (particularly in novel environments). This talk aims to explore this idea, suggest a potential pipeline for end-to-end machine learning for robust robot control, and promote discussion and feedback around this solution's plausibility from people who know much more about machine learning than I do.

Sharon Temtsin

Title: A Bona fide Turing Test

Abstract: Conversational agents based on Large Language Models (LLM) have been rapidly adopted by users since the public release of ChatGPT. These systems have demonstrated exceptional performance that has never been witnessed before in their capability to generate human-like texts and responses. They have excelled not only in their conversational abilities but also in dealing with tasks such as professional exams, and complex problem-solving. Such performance fosters the expectation that we are on the brink of creating an intelligent machine. However, measuring intelligence remains elusive without a proper test. Alan Turing proposed a test for machine intelligence based on imitation and unconstrained conversations between a machine and a human. To the best of our knowledge, no one has ever conducted Turing's test as Turing prescribed, even though the Turing Test has been a bone of contention for more than seventy years. Conducting a bona fide Turing Test would not only significantly contribute to the advancement of research in machine intelligence evaluation but also hold the potential to advance AI researchers in their ultimate quest of developing an intelligent machine.

Sreedevi Sandara Iyer

Title: Empathy measure for software professionals

Context: Empathy among software professionals can reduce work-related stress and create a supportive and conducive working environment that supports the success of software projects. Since empathy is crucial for software professionals, it is important to assess this skill to identify their levels of competence in empathy as a way to propose developmental needs as necessary.

Objective: To develop a validated measurement instrument that can assess software professionals' levels of empathy

Method: Items for the Empathy Measure for Software Professionals (EMSP) were developed through a literature review of existing measures of empathy. These items were then subject to content analysis by academic experts to establish the item's relevance, clarity, and singularity. Currently, this empathy measure is being piloted on second-year Bachelor's with Honours software engineering students to

identify any issues with the structure and content of the measure. The data collected from this pilot study will be used to conduct an initial exploratory factor analysis to remove any unwanted items and identify the factor structure. This modified empathy measure will be further subject to yet another exploratory factor analysis to hypothesise the factor structure, which will established through a confirmatory factor analysis. This study will also perform tests of validity such as predictive, convergent, and discriminant validity to establish the extent to which this instrument is valid.

Findings: The findings of the study are expected to contribute towards a validated and reliable empathy measure for software professionals to assess their empathy towards their team members.

Sujan Warnakulasooriya

Title: Novel method for optimising initial drone flocking

Abstract: Unmanned Aerial Vehicles (UAVs) which are also known as drones can significantly improve their performance when working as a flock. One key area of research under drone flocking is the formation of the initial flock. The focus of this proposed research would be to investigate how to optimise the initial formation of a flock of N drones, with the goal of minimising drone travel distance, energy consumption, and initial formation time. Two special cases will also be investigated: analysing the delayed arrival of a single drone in a fixed formation and reducing the initial transient response of a flock of drones. This research aims to improve the performance of drone flocks.

Tim Rensen

Title: Photogrammetry and machine learning for surveying scallops.

Abstract: Scallop fishery SCA7 at the top of the South Island was closed due to stock collapse in 2017, halting commercial and recreational fishing of hundreds of tonnes of scallops, and rahui or restricted access imposed on customary fishing by local Māori. Many factors contributed to this stock collapse, including overfishing, sediment and nutrient runoff from land use, and natural events such as once-in-a-hundred-year floods turning productive seabed into sediment muck meters deep, suffocating and starving seabed communities. Scallop dredging involves raking up scallops off shallow seabed where they reside, upturning delicate seafloor, tearing apart microalgal films, resuspending sediment, and permanently changing the topology of the seabed.

For commercial scallop fishing to be more sustainable we need better harvesting and monitoring methods. Autonomous underwater vehicles (or AUVs) make if feasible to scan large areas of seabed consistently. The image data can be processed into 3D reconstructions of habitats, allowing marine scientists to analyse entire benthic habitats with a single 3D model or mosaic image as opposed to frame-by-frame in a video.

I propose an annotation method where scientists can label a reconstructed top-down image of a seabed environment, and these labels propagated to all of the individual image frames. This method requires labelling of unique scallops only once, avoiding having to label the same scallop in multiple frames. Transfer learning will be used to reduce the required amount of data in training a CNN for pixel-wise identification of scallops.

How this will help in monitoring and harvesting of scallops is in development of more efficient methods for analysing the scallop stocks and habitats from AUV video. This same robot can run scallop classification live to locate harvestable scallops and selectively pick them up, eliminating bycatch and minimising

habitat damage. The result would be more sustainable scallop fisheries, and increased environmental monitoring, enhancing kaitiakitanga or our stewardship of earth.

Zhouyu Qu

Title: Wireless path planning algorithms on a realistic drone road system

Abstract: Copter-type UAVs (unmanned aerial vehicles) or drones are expected to become more and more popular for deliveries of small goods in urban areas. One strategy to reduce the risks of drone collisions is to constrain their movements to a drone road system as far as possible. For reasons of scalability, we assume that path-planning decisions for drones are not made centrally but rather autonomously by each individual drone, based solely on position/speed/heading information received from other drones through WiFi-based communications. We present a system model for moving drones along a straight road segment or tube, in which the tube is partitioned into lanes. We furthermore present a cost-based algorithm by which drones make lane-switching decisions, and evaluate the performance of differently parameterized versions of this algorithm, highlighting some of the involved tradeoffs

The organisers are grateful for the support of the following sponsors:

Principal Sponsor

Verizon Connect builds solutions and services that put innovation, automation and connected data to work for customers and help them be safer, more efficient and more productive.

With more than 3,500 dedicated employees in 15 countries, we deliver leading mobile technology platforms and solutions.

Verizon Connect offers web-based mobile workforce and fleet management software including GPS tracking, route planning, job management, navigation and compliance.

Commercial fleets and mobile workforces of all sizes use our software to improve the effectiveness and efficiency of their operations – and reduce their carbon footprint.

Verizon Connect began in part a Canterbury success story. Ralph Mason, a founder of one of the companies that came together as Verizon Connect, hails from Christchurch and set up a research and development centre here 14 years ago.

Much of the original technology was developed in Christchurch and has been instrumental in shaping this billion-dollar industry. As a result, Verizon Connect is considered a technology leader and pioneer in the telematics industry.

Verizon Connect has offices in 15 countries around the world including Dublin, London, Sydney, Los Angeles, Chicago, Austin with staff numbers over 3,500.

The company has a large research and development centre based in Christchurch.

Supporting the development of its SaaS software are a range of development and test automation roles requiring capability in C#, JavaScript, C++, REST, SOA, UX, UI, systems administration and database management.

Verizon Connect regularly hires for its Christchurch office with at least 10 graduates each year and five to 10 interns working year-round. Many interns stay on as permanent employees after graduation.

The company is actively looking for software developers and engineers as well as business analysts, support engineers, software test engineers and software development managers.

We look for staff who can embrace change and new ways of thinking – and move quickly

We're a diverse mix of professionals who thrive in an atmosphere that promotes fresh ideas and challenges traditional systems.

Our culture is one that values open communication, and is built on creativity and excellence

Contact us: verizonconnect.com/company/careers

Tait customers protect communities, power cities, move people, harness resources and save lives all over the world. We create and support their critical communications.

From our strong position as leaders in radio communication technology, we work hard to gain a deep understanding of the issues, problems, and day-to-day working environments our customers' experience. That is how we deliver robust, fit-for-purpose products, exceptional customer service, and world class communication system performance.

Our LMR products and systems are designed and built by our people. We stand by their quality, integrating, testing and perfecting everything we sell. Our specialties include P25 (Phase 1 and Phase 2), DMR (Tier 2 and Tier 3), MPT-1327, Analog Conventional, and software to manage and monitor those radio networks.

We also recognize the increasing importance of business systems like dispatch, AVL, LTE, cellular, voice recorders, SCADA, and many more to come. We integrate these diverse technologies to deliver stronger, simpler, and smarter solutions. Through these efforts, we're redefining critical communications.

Allied Telesis: For nearly 30 years, Allied Telesis has been delivering reliable, intelligent connectivity for everything from enterprise organizations to complex, critical infrastructure projects around the globe.

Allied Telesis is recognized for innovating the way in which services and applications are delivered and managed, resulting in increased value and lower operating costs.

Originally part of the DSIR, Allied Telesis Labs - based in Christchurch, New Zealand - joined the international Allied Telesis Group in 1999, and is now the largest research and development centre for the group.

As one of the world's leading producers of computer networking equipment, the group employs more than 3000 people worldwide.

Allied Telesis smart technologies, such as Allied Telesis Management Framework™ (AMF) and Enterprise SDN, drive network evolution, and deliver efficient and secure solutions for people, organizations, and "things".

More than one million customers worldwide have chosen our technology for their networks.

The success of Allied Telesis Labs is built on the skills of our talented employees, who enable the company to compete on an equal footing with the world's largest communications and networking equipment manufacturers.

www.alliedtelesis.co.nz

Founded in 1993 and 100% New Zealand owned, PB Tech is NZ's largest computing hardware reseller with stores + service centres throughout the country and a hardworking team of over 800 staff.

In addition to our retail and service locations, PB Tech also has dedicated Business, Education, Wholesale, Government, Health and Service divisions operating from our 10,000+ square metre head office and distribution centre in Manukau, Auckland

Over the last 30+ years, we've worked hard to build a trusted reputation as a leading supplier of technology products and services for home, business and the public sector, and count many of New Zealand's largest companies, government departments and educational institutes as our customers.

Much of PB's success is the result of an unwavering commitment to helping our customers save money by offering New Zealand's lowest prices on a huge range of computing and technology hardware. Our obsession with being NZ's best value shopping destination for all things tech is complimented by our unmatched selection of in-stock and ready-to-ship products, convenient shopping locations in-store and online, and passionate and knowledgeable staff.