UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Software Engineering Final Year Projects
-
Showcase programme

2023

Novotel, Cathedral Square
Ōtautahi Christchurch

Computer Science & Software Engineering

Software Engineering Final Year Projects

Showcase

Novotel Christchurch Cathedral Square, Christchurch (New Zealand)

12^{th} of October, 2023

Contents

Foreword

The Final Year Project, *aka SENG402 Software Engineering Research Project*, is a required course for all graduating software engineering students. This research project provides the opportunity for students to demonstrate they are prepared to work as professional software engineers. Students build on the skills learned through projects they have undertaken in earlier years of the programme, in particular a single-semester project in second year (*SENG202 Software Engineering Project Workshop*) and a whole-year project at third year (*SENG302 Software Engineering Group Project*). SENG402 projects offer a specific research-focused venue for students to explore advanced methods or late breaking results into new problems.

The department offers students the opportunity to engage with industry on company-sponsored projects. These opportunities complement the range of projects offered each year by academics in the department. This year, more than 30 industry and academia projects were offered to the software engineering students with half of the students engaging with the industry where cutting edge research expertise can be applied to an industrial context.

The final year presentations are part of the formal assessment for students. More importantly, this showcase day provides us all with the opportunity to recognize and celebrate the achievements of our final year software engineering students as they prepare to graduate. This is the moment to thank all our students for their commitment and contribution to the department, College, University as well as the industry.

Fabian Gilson & Etienne Borde, SENG402 course coordinators

On behalf of the Computer Science and Software Engineering Department

Summarised schedule

	Room:Waimakariri 1	Room:Waimakariri 2
9:00-9:15	**Welcome Words**	
9:15-10:30	Education chair: Fabian Gilson	Geo-localisation and classification chair: Etienne Borde
10:30-11:00	**Morning Tea**	
11:00-12:15	Software Engineering chair: Tanja Mitrovic	Automated machinery chair: Chenyi Zhang
12:15-13:15	**Lunch**	
13:15-14:30	Agri-tech chair: Miguel Morales-Trujillo	Environmnent chair: Trevor Nesbit
14:30-15:00	**Afternoon tea**	
15:00-16:15	Health and Care chair: Andreas Willig	Varia chair: Marina Filipovic
16:15-16:30	**Closing Words**	
16:30-17:30	**Social Event**	

Presentation coloured badges

I Industry project R Research project P Personal project

Track: Shafi Goldwasser (Turing award 2012)

Room: Waimakariri 1

9:00-9:15		**Welcome Words**	
9:15-9:30	R	**April Clarke** *University of Canterbury*	Explainable AI in an Intelligent Tutoring System
9:30-9:45	R	**George Carr-Smith** *University of Canterbury*	Gamification in Educational and Industrial Contexts: Play and Learn
9:45-10:00	R	**George Hampton** *University of Canterbury*	CodeWOF Exercise Creation Interface
10:00-10:15	R	**Samuel McMillan** *University Of Canterbury*	Unlocking SQL Mastery: A Gamified Approach to Enhanced Learning with QueryCompetition
10:15-10:30	R	**Sami Elmadani** *University of Canterbury*	EER Tutor – Self Regulated Learning
10:30-11:00		**Morning Tea**	
11:00-11:15	R	**Darryl Anne Alang** *University of Canterbury*	State of DevSecOps Practice in Aotearoa New Zealand
11:15-11:30	I	**Luke Garside** *University of Canterbury, Department of Physics*	Barracuda: A language for executing arbitrary GPU code at runtime
11:30-11:45	R	**Danish Khursheed Jahangir** *University of Canterbury*	Coding with Confidence: The Role of Tools in Code Review Excellence
11:45-12:00	R	**Saskia van der Peet** *University of Canterbury*	Social sustainability in the New Zealand software industry
12:00-12:15	I	**Sahil Negi** *Ministry of Education*	Connecting physically remote families through stories
12:15-13:15		**Lunch**	
13:15-13:30	I	**Fa Wren Chong** *Ministry of Business, Innovation and Employment*	Texture Boundary Segmentation for Pinus Radiata Branch Cut-Point Detection
13:30-13:45	R	**James Claridge** *MBIE*	Tree branch detection for Drones
13:45-14:15	I	**Arish Myckel Abalos, Andrew Cook** *Onside Limited*	Agricultural Land Identification and Classification utilising Deep Learning
14:30-15:00		**Afternoon tea**	
15:00-15:15	I	**Cody Larsen** *E-Quip (Hannes de Bruin)*	Constant Connect Memory Supportive Technology
15:15-15:30	R	**Dillon Pike** *University of Canterbury*	Real-time Hip Inserter Angle Tracking
15:30-15:45	P	**Yiyang (Jessie) Yu**	Deep Learning for Diabetic Retinopathy
15:45-16:00	P	**Rachel Hodgson**	Patient Deterioration Detection

16:00-16:15	R	**Samuel Clark** *University of Canterbury, Department of Computer Science and Software Engineering*	Texture Feature Analysis of Breast Cancer Regions in Whole Slide Images
16:15-16:30		**Closing Words**	
16:30-17:30		**Social Event**	

Track: Liskov Barbara (Turing award 2008)

Room: Waimakariri 2

9:00-9:15		**Welcome Words**	
9:15-9:30	I	**Billy Sandri** *Trimble*	Fusion of GNSS with Mobile SLAM
9:30-9:45	R	**Amy Sloane** *University of Canterbury*	Mobility Data Visualization - A Tool for Detecting Suspicious Geospatial Patterns
9:45-10:00	R	**Jamie Thomas** *University of Canterbury, FOAM*	Proof-of-location testing with a FOAM trust zone on UC campus
10:00-10:15	I	**Jack McCorkindale** *Seequent*	Mixed Reality Collaboration in Geoscience
10:15-10:30	I	**Prableen Oberoi** *Trimble*	Augmented Reality GIS Assets
10:30-11:00		**Morning Tea**	
11:00-11:15	I	**Lachlan Alsop** *Zealandia Systems, University of Canterbury*	CNC Milling Chatter Detection Using a CNN AI
11:15-11:30	R	**Hugo Reeves** *Univeristy of Canterbury*	Insights Into Power Line Detection and Distance Estimation Using Stereo Cameras
11:30-11:45	I	**Bede Skinner-Vennell** *Unison*	Conductor Condition Recognition Using Autonomous Drones
11:45-12:15	R	**Euan Morgan, Frederik Markwell** *University of Canterbury*	Smooth Drone Formation Flight
12:15-13:15		**Lunch**	
13:15-13:30	I	**Jacques Terblanche** *Antarctica NZ*	Scott Base Seal Census using Machine Learning
13:15-13:30	R	**Karl Moore** *University of Canterbury Biosecurity Innovations*	Mobility Data for Biosecurity
13:30-13:45	R	**Jonathan Tomlinson** *Simon Kingham, University of Canterbury*	Using Mobile Location Data to Inform Climate-Resilient Transport Interventions
13:45-14:00	I	**Lane Edwards-Brown** *AgResearch*	Port Identifier for Pests (PIP)
14:00-14:15	R	**Shaylin Simadari** *University of Canterbury*	Efficient Annotation of Sparse Vocalisations for Bioacoustic Monitoring
14:30-15:00		**Afternoon tea**	
15:00-15:15	R	**Oliver Johnson** *CSSE*	Sentiment analysis for licensing applications

15:15-15:30	I	**Max Bastida** *Eagle Technology*	Digital Twin Builder
15:30-15:45	I	**Aidan Campbell** *Institute of Environmental Science and Research*	Bloodstain Pattern Classification using Deep Learning
15:45-16:00	I	**Moses Wescombe** *The Institute of Environmental Science and Research (ESR)*	Laboratory automation at ESR
16:15-16:30		**Closing Words**	
16:30-17:30		**Social Event**	

Track: Shafi Goldwasser (Turing award 2012)

Session: Education

(time: 9:15-10:30)

Chair: Fabian Gilson

Explainable AI in an Intelligent Tutoring System

R

April Clarke University of Canterbury

Supervisor: Tanja Mitrović

(Waimakariri 1, 9:15-9:30)

Intelligent Tutoring Systems (ITSs) are educational systems that provide adaptive learning environments based on students' knowledge, preferences, and needs. SQL-Tutor, an ITS that teaches SQL select queries, improves learning outcomes for students, particularly when they work on problems recommended by the system. However, a key challenge lies in students' varying willingness to accept the system's recommendations. This may be due to a lack of clarity in the reasoning behind these recommendations. To address this challenge, we implemented explanations for the system's problem selection recommendation. Then, we studied students' responses to the explanations.

Our research focuses on the following questions:

- How does students' engagement with explanations affect their interactions with an ITS?
- What questions do students want answered by explanations in an ITS?
- How do individual personality traits influence students' interactions with explanations in an ITS?

Prior research shows that these explanations typically increase students' trust, and benefit less conscientious students the most. Our findings will help us gain more insight into how we can use explanations to improve students' learning experiences, and which explanations students value the most.

Gamification in Educational and Industrial Contexts: Play and Learn

R

George Carr-Smith University of Canterbury

Supervisor: Miguel Morales Trujillo (Waimakariri 1, 9:30-9:45)

Research Problem: The purpose of this study is to investigate the potential increase in student engagement, motivation, and satisfaction resulting from the use of a gamified approach to teaching risk management. This study implements a digital version of Pictionary called "Risktionary" to educate students about risk management. A selected player illustrates a risk on a digital whiteboard, while the other students attempt to identify the risk and vote on its probability and severity.

Research Question: "How do collaborative game elements in a game-based solution enhance student motivation, engagement, and satisfaction in a course at a tertiary education institution?"

Situating the Case: Similar studies predominantly implemented the physical format for Pictionary, while some studies adapted it to a digital game. These studies revealed that students were enthusiastic about using Pictionary as a learning tool and expressed a desire for more game-based learning in their education.

Methodology: This study employs a qualitative analysis of surveys from 10 participants to gauge the levels of engagement, motivation, and satisfaction experienced while playing Risktionary. Additionally, the study quantitatively analyses the game's effectiveness by comparing tests administered to students before and after gameplay. The primary design cycle for this study is an action feedback loop, which aims to iteratively enhance the game based on feedback gathered from student surveys.

Conclusions: Utilising Risktionary to teach software risk management moderately enhanced students' learning and consistently maintained their engagement throughout the game.

CodeWOF Exercise Creation Interface

George Hampton University of Canterbury

Supervisor: Tim Bell

(Waimakariri 1, 9:45-10:00)

This project aims to enhance the CodeWOF website (which allows teachers to combat skill fade by completing regular exercises). The prior way to add questions relied on manually creating YAML files, and as such was prone to human error. Additionally, variations of a question were added by creating a full identical question and changing it slightly. Therefore, a more human-friendly interface which better supports question variation was required.

As part of developing the interface, a study was conducted into whether being able to create questions would benefit teacher users. This was used to determine whether the interface should be available to teachers, or only administrator users. The developed interface allowed users to create new questions and optionally define randomisation to vary questions.

Unlocking SQL Mastery: A Gamified Approach to Enhanced Learning with QueryCompetition

R

Samuel McMillan University Of Canterbury

Supervisor: Miguel Morales-Trujillo

(Waimakariri 1, 10:00-10:15)

In the domain of education, gamification, or the introduction of game elements in non-gaming environments, has shown to hold transformative potential. This final-year project delves into the nuanced effects of gamification, with a particular emphasis on competitive and non-competitive elements, on students mastering the Structured Query Language (SQL). At the heart of this study is the web application, QueryCompetition (QC), a tool that offers an environment where students hone their SQL skills through interactive challenges and are incentivized via a point-based system, gauging factors like time efficiency and query accuracy. This research endeavors to amplify QC's efficacy, subsequently conducting empirical studies to derive insights into gamification's tangible impact on SQL learning outcomes. Preliminary investigations suggest avenues for augmenting student retention. The project's trajectory entails a strategic blend of QC refinements, followed by a more rigorous empirical assessment, all pointing towards a constructive and promising outcome for SQL education.

EER Tutor – Self Regulated Learning

Sami Elmadani University of Canterbury

Supervisor: Tanja Mitrovic

(Waimakariri 1, 10:15-10:30)

This research project aims to integrate Self Regulated Learning (SRL), guided by the Zimmerman model, into the Intelligent Tutoring System known as EER-Tutor. EER-Tutor is designed to offer personalized support to students during database design. SRL is a dynamic learning process through which students actively manage their own learning using cognitive, motivational, and behavioral strategies. The Zimmerman model, a central framework in this study, defines SRL as a cyclical process comprising three distinct phases: forethought, performance, and self-reflection. Within these phases, students engage in task analysis and goal setting, employ strategy selection, self-monitoring, and apply chosen strategies, and they perform self-evaluation and adaptations for future learning.

The objective to enhance students' learning and SRL is proposed with the implementation of three key modules within EER-Tutor. Goal Setting: Students can set specific, challenging, and attainable session goals, fostering self-regulation skills. Pre-Post Self-Assessment: Before attempting problems, students analyze task requirements and strategies, enhancing their strategic planning abilities. Afterward, they reflect on their performance. Problem Selection: Students choose their next problem using methods that suite them, promoting self-reflection and planning based on past experiences.

The integration of these SRL modules aims to enhance student performance and offers practical insights with potential benefits for educators and instructional designers.

Session: Software Engineering

(time: 11:00-12:15)

Chair: Tanja Mitrovic

State of DevSecOps Practice in Aotearoa New Zealand

R

Darryl Anne Alang University of Canterbury

Supervisor: Fabian Gilson

(Waimakariri 1, 11:00-11:15)

Security evaluations are tedious in nature. Due to the constant reduction in the time-to-market of software products, rapid software development and deployment led to the negligence of security requirements. This is evident with the recent increase in security attacks against organizations across various industries, including New Zealand. Recognizing the need for speed and security, the concept of DevSecOps emerged.

This study aims to determine the current state of practice of DevSecOps in New Zealand, and the benefits gained and the challenges an organization face when implementing such process.

To achieve this, we conducted a literature review where we reviewed relevant articles from Scopus. We found that despite the push for "security to the left" mindset, most artefacts emphasized the importance of continuous testing and monitoring. Our review also revealed that the main challenges faced by practitioners are speed, lack of knowledge about the common vulnerabilities and tools, and change in culture.

Moreover, we conducted a survey with professionals in the software industry. The results indicated that about 40% of the respondents do not receive security training. However, most of them consider security risks when writing code and functional requirements. The responses regarding the security tools used were diverse, while the majority indicated that linters and static analysis tools were used. Additionally, they consider the lack of experience or knowledge about security and the time-to-market pressure to be the major challenges when integrating security practices.

Barracuda: A language for executing arbitrary GPU code at runtime

I

Luke Garside University of Canterbury, Department of Physics

Supervisors: Fabian Gilson, Phillip Duncan

(Waimakariri 1, 11:15-11:30)

Physics simulations often need to execute arbitrary code at runtime. Programming languages often have features that allow users to do this. However, no such function exists for GPU languages.

The Barracuda library has been created to solve this problem. It works with CUDA, a GPU programming language created by Nvidia. Barracuda takes in code in the form of a list of primitive instructions, allowing for arbitrary code to be executed at runtime. However, this is a difficult format to code in for developers.

A compiler for Barracuda has been created so developers can write in a higher-level language that compiles to a list of instructions for the Barracuda library. A compiler for Barracuda was written in 2022. However, it contained bugs, was missing many features, and produced slow code.

This project aims to improve the usability of the Barracuda compiler. To do this, we have written tests to discover bugs and increase confidence in the compiler. We have fixed many bugs in the compiler, and it is now working. We have extended the compiler with pointers, arrays, and a type system. Arrays were needed as without them, programs involving lists of data were very hard to write. Pointers were needed as arrays are a kind of pointer. A type system was needed to implement complex types. For example, an array of pointers to arrays. Implementing a type system involved creating a new step for the compiler, a semantic analyser. We have also investigated potential optimizations for the compiler.

The project is useful for any developers who want to allow arbitrary code to be executed at runtime on GPUs. For example, the Physics Department at the University of Canterbury has created a program that simulates an MRI machine that uses Barracuda.

Coding with Confidence: The Role of Tools in Code Review Excellence

R

Danish Khursheed Jahangir University of Canterbury

Supervisor: Fabian Gilson

(Waimakariri 1, 11:30-11:45)

Code reviews are a process in software development where developers examine and evaluate each other's code changes. This helps them enforce coding standards, detect overlooked bugs, ensure code consistency, and more. However, when dealing with a large number of code changes, it can become challenging to thoroughly review every aspect. Additionally, individual developers may have their own unique criteria for conducting reviews, which may affect the code review process.

This presents a significant challenge as it results in inconsistent code reviews and can consume considerable time, increasing the likelihood of human errors and overlooked issues during the review process.

To address these challenges, our research project is dedicated to identifying and implementing best practices and tools that enhance the efficiency and effectiveness of code reviews. To this end, we first conducted a comprehensive literature survey to compile current best practices and tools widely used in the software development industry. Next, we engaged in focused discussions with 3rd-year Software Engineering students. We encouraged these students to utilize specific tools such as SonarQube, merge request checklists, CI/CP pipelines, among others, during their code review processes. These discussions aimed to gain insights into their experiences, methodologies, and the impact of these tools on their code review practices.

By combining industry best practices with real-world insights from these student discussions, our research strives to refine and optimize the code review process, and find how the usage of code review tools affects the code review process in software development projects.

Social sustainability in the New Zealand software industry

Saskia van der Peet University of Canterbury

Supervisors: Miguel Morales-Trujillo, Ismael Caballero Muñoz-Reja (Waimakariri 1, 11:45-12:00)

In 2015, all United Nations member states, including New Zealand, committed to the 2030 Agenda for Sustainable Development, and its 17 Sustainable Development Goals. With software playing such a crucial role in society, it is also important that it contributes to our commitment to the 2030 Agenda for Sustainable Development.

The objective of our study, was to understand the current practices and challenges of social sustainability, within the New Zealand software industry. We prioritised social sustainability, which focuses on community impact, because it has the most influence on the Sustainable Development Goals.

To achieve this objective, we conducted a systematic literature review, survey and interviews with New Zealand software industry professionals.

Preliminary results indicate that for some organisations, sustainability is part of their mission. Others have processes in place to ensure sustainable decision making and product development, or support the community through initiatives, such as fundraising, or volunteering.

Connecting physically remote families through stories

I

Sahil Negi Ministry of Education

Supervisor: Fabian Gilson

(Waimakariri 1, 12:00-12:15)

This project addresses the challenge of geographic separation within families, which hinders shared storytelling experiences, particularly for young children and their extended loved ones. Our innovative app seamlessly combines physical books with e-books to bridge geographical gaps, strengthening family connections.

Our objectives include enhancing literacy and fostering stronger bonds among physically remote families. The app empowers users to access the overdrive library page, upload PDFs, scan books, view uploaded PDF books, and engage in video calling features between caretaker and child users. By providing a user-friendly and engaging platform, our ultimate goal is to deliver a reliable app that connects physically remote families, promotes literacy, and enhances family connections.

Session: Agri-Tech

(time: 13:15-14:30)

Chair: Miguel Morales-Trujillo

Texture Boundary Segmentation for Pinus Radiata Branch Cut-Point Detection

I

Fa Wren Chong Ministry of Business, Innovation and Employment

Supervisors: Prof Richard Green, Sam Schofield

(Waimakariri 1, 13:15-13:30)

The Pinus Radiata species of pine tree is pruned to produce valuable clear wood. Currently, this task is performed manually, which requires an arborist to traverse dense gorse and blackberry plants to reach the tree. Additionally, the use of ladders to prune branches at height poses a hazard to the arborist.

Unmanned Aerial Vehicles (UAVs, or drones) are well suited to the task of pruning Pinus Radiata trees. A crucial step in automatic pruning is the detection of the precise cut-point on the branch, as incorrect pruning caused by inaccurate detection may damage the trunk and produce undesirable discolouration in the resulting wood. For this purpose, a dataset was produced by augmenting and annotating photos of Pinus Radiata tree branches captured on location and supplemented with images of a model tree, for a total of 158 images. This dataset was used to train the R-CNN deep neural network to detect the cut-point region. The model's detection was evaluated at distances ranging from 0.5 to 4.5 metres from the branch in a controlled environment. Detection was achieved at a distance of approximately 1 metre from the branch.

Tree branch detection for Drones

James Claridge MBIE

Supervisor: Prof. Richard Green (Waimakariri 1, 13:30-13:45)

This project aims to develop software to aid in the autonomous pruning of tree branches by drones. The UC Vision group is developing a system to allow pruning of tree branches autonomously in the foresty industry. This system calls for the ability to detect branches for which a deep learning model is being developed. This project is concerned with developing this deep learning model.

The proposed solution uses a Segformer model trained on branch images to detect branch pixels within a given image. The model was pretrained on a large image dataset and then finetuned on images of branches.

Agricultural Land Identification and Classification utilising Deep Learning

I

Arish Myckel Abalos, Andrew Cook Onside Limited

Supervisor: Mukundan Ramakrishnan

(Waimakariri 1, 13:45-14:15)

Onside Limited is currently enhancing their data-driven biosecurity model, Onside Intelligence (OSI). By using the data based on what they such as contractor movements between properties and what is on properties, OSI aims to deliver a streamlined management method of biosecurity threats when they occur, rather than the current model which is checking properties randomly. By identifying traces of biosecurity threats, officials can determine the impact, the spread, and their next steps of mitigation from the data that OSI provides. Currently, the process for adding more resolution to what is on properties is by using "blocks", which are sub-regions within the property usually designated to different crop types or agricultural land use. Our project aims to assist OSI by automating this process as much as possible, making the onboarding process for customers easier and allowing for more data to be available for OSI.

This project is split into two subprojects. The first is identification of boundaries of blocks, and the second is the classification of what is inside of those boundaries. Both sub-projects utilise Computer Vision techniques and Deep Learning using LINZ's Basemap data. LINZ's Basemap data provides satellite and aerial photography anywhere between 10 meters per pixel, to 0.025 meters per pixel. This project went in steps of data sourcing, image processing, data labelling, model training, and model evaluation to achieve automatic mapping and classification of blocks based on the LINZ's Basemap dataset.

Session: Health and Care

(time: 15:00-16:15)

Chair: Andreas Willig

Constant Connect Memory Supportive Technology

I

Cody Larsen E-Quip (Hannes de Bruin)

Supervisor: Fabian Gilson

(Waimakariri 1, 15:00-15:15)

This project addresses the challenges faced by dementia patients who heavily rely on their family members and caregivers, wanting to talk to their families multiple times per day asking the same repetitive questions. This can place a heavy burden on families of dementia patients and cause a lot of stress and frustration.

"Constant Connect" is an innovative assistive technology that aids dementia patients by simulating real conversations with family members. It employs advanced voice recognition technologies and pre-recorded responses from family members, allowing dementia patients to experience the comforting presence and recognition of their loved ones while also reducing the burden on those family members being overwhelmed by repetitive questions.

The solution also aids dementia patients by modeling the look of the technology to resemble an old telephone to give the patient a sense of familiarity. The opportunity to reminisce on family and personal memories, as well as messages from loved ones, is also available to dementia patients through easily accessible photo and video galleries. Additionally, scheduled and intermittent reminders are included within the technology to aid with essential patient care, such as reminders to drink water, take medicine, and more.

Real-time Hip Inserter Angle Tracking

Dillon Pike University of Canterbury

Supervisor: Richard Green

(Waimakariri 1, 15:15-15:30)

Hip replacement surgeries are becoming increasingly common, especially for younger patients. The success of the surgery hinges on the placement of a semi-spherical cup within the pelvis to create an artificial ball-and-socket joint. Inaccurate cup insertion can lead to post-operative issues, such as limps.

This project aims to improve the surgery's reliability without requiring any modifications to the hip inserter. This is done by providing surgeons with real-time pitch and roll angles of the hip inserter, enabling them to achieve more accurate cup placements.

The approach uses a stereo camera to capture left and right video feeds. A Convolutional Neural Network (CNN) is employed to detect feature points along the hip inserter in both video feeds. The points in each video feed are combined to form 3D points, which are used calculate pitch and roll angles. Initial testing demonstrates that a well trained CNN can achieve real-time angles with an average error of one or two degrees.

Deep Learning for Diabetic Retinopathy

P

Yiyang (Jessie) Yu University of Canterbury

Supervisor: Andrew Bainbridge-Smith

(Waimakariri 1, 15:30-15:45)

Diabetic retinopathy (DR) is a leading global cause of vision impairment that can be prevented by early detection, periodic retinal screenings, and medical treatments. By 2040, over 600 million people with diabetes will need regular eye checks. With diabetes on the rise, it's crucial to ensure that the future of our health system can provide prompt access to DR screenings. New Zealand's current manual screening process is slow and relies on professionals to read images accurately, causing a stressful bottleneck to the number of patients the system needs to expand to manage.

Automating screenings using predictive technology can assist in detecting and grading the severity of the disease. While previous studies have explored this, there is still room for improvement with a deep learning model tailored for New Zealand. To address this, a deep learning model was trained on New Zealand-specific retinal images, aiding medical professionals in accurately classifying DR stages, improving healthcare access, and ensuring timely intervention for affected individuals.

Patient Deterioration Detection

P

Rachel Hodgson University of Canterbury

Supervisor: James Atlas

(Waimakariri 1, 15:45-16:00)

This project addresses the challenge of detecting rapid health deterioration in patients within emergency department waiting rooms. Current triage systems in New Zealand assign patients a priority based on symptoms, vital signs, and temperature, however, there is a lack of reassessment between triage and seeing the doctor, which can lead to undetected deterioration.

Current deterioration detection systems implemented on wards use traditional contact sensors including heart rate monitors, blood pressure cuffs, contact thermometers, etc. While these contact sensors are accurate, they are impractical for monitoring every patient in a waiting room and have negative physical impacts for continuous monitoring, such as skin irritation and patient discomfort. A camera-based approach, such as this project, offers a non-contact and cost-effective solution, allowing a single system to monitor multiple patients.

The developed solution uses a Logitech WebCam and a FlirOne thermal camera to detect when a patient stops moving or when their body temperature is too high. The solution prints a notification to a console to alert the user that a deteriorating patient has been discovered. In an industrial setting, this would equate to a nurse or doctor being notified that a patient was experiencing deterioration, allowing medical intervention to take place, improving the health outcomes for patients.

The motion detection system demonstrated an accuracy rate of 90%, with a precision of 100% and a recall of 80%. The risk profile identified the most critical issue being the occurrence of false positive results, where the system did not detect a lack of motion, failing to inform the absence of movement. This anomaly manifested 0 times during testing, yielding a false positive rate of 0%.

Alongside this, the fever detection system exhibited an accuracy rate of 92%, a precision of 83%, and a recall of 100%. Similar to the motion detection scenario, the primary concern was the generation of false positive results, where the system failed to identify a fever that was present. This issue was encountered four times during testing, resulting in a false positive rate of 13%.

However, there are limitations. While the motion detection can monitor multiple people at once, the heavy computation time of the object detection algorithm creates difficulties in identifying small movements, resulting in 20% false negatives. This could be mitigated by exploring ways of making the object detection more efficient. The thermal camera monitors the whole environment for fevers, resulting in false negatives for non-human objects that have a temperature of over 38.3°C/100.94°F. The thermal camera also only measures surface temperature, which may not be entirely accurate for fever detection. Additionally, the system is not currently a live feed due to equipment limitations.

Future research to improve the motion detection could explore more efficient person detection algorithms to reduce false negatives. Future research for thermal analysis could investigate the use of person detection to disregard hot objects in frame and the implementation of USB thermal cameras for live analysis.

Texture Feature Analysis of Breast Cancer Regions in Whole Slide Images

R

Samuel Clark University of Canterbury, Department of Computer Science and Software Engineering

Supervisor: Prof. Ramakrishnan Mukundan (Waimakariri 1, 16:00-16:15)

Breast cancer is a prevalent and life-threatening disease, affecting millions of individuals worldwide. Early detection and accurate diagnosis are critical for improving treatment outcomes and reducing the mortality rate. Currently, tissue samples of breast cancer regions are evaluated manually by pathologists to classify cancer regions and recommend treatment plans.

Pathologists diagnose and evaluate breast cancer through microscopy of tissue samples. However, manual analysis of histopathological samples is a laborious task and yields variable results, creating an opportunity for computational pathology to assist in breast cancer biopsies. Computational pathology uses whole slide images of histopathological samples and analyzes them to extract characteristics of areas of interest. This simplifies biopsies for pathologists, allowing them to focus their attention on the regions of interest.

Our study aims to create a machine learning model to accurately segment cancerous regions from whole-slide images. We are currently analyzing the model to determine which texture features are most relevant for identifying cancer regions. This is achieved by extracting a set of 964 texture features, utilizing feature extraction algorithms such as Local Binary Pattern (LBP), Grey Level Co-occurrence Matrices (GLCM), and statistical measures. These features are then used to train a random forest classifier.

Track: Liskov Barbara (Turing award 2008)

Session: Geo-localisation and classification (time: 9:15-10:30)

Chair: Etienne Borde

Fusion of GNSS with Mobile SLAM

Billy Sandri Trimble

Supervisor: Fabian Gilson (Waimakariri 2, 9:15-9:30)

TerraFlex, a mobile application from Trimble's Geospatial department, allows customers to collect information about their assets. An asset is any object of interest to the customer such as a tree, rubbish bin or the corner of a building. TerraFlex provides high accuracy GNSS positioning via a DA2 receiver and Trimble's real time atmospheric correction service. The accuracy of the GNSS position degrades under tree cover or in between buildings, where there are less satellites with line of site to the receiver.

We have built a **Fusion Engine** which utilises the camera on a mobile phone to assist the position of the GNSS receiver. We are able to estimate the relative motion of the mobile phone by tracking how far it has moved between frames of the camera. This estimated motion can be used to predict the global location of the receiver, based on its last accurate position.

When a TerraFlex user has camera assistance enabled, they will be provided with a **fused position**. This is an abstraction layer for the receiver and camera positions. When GNSS accuracy drops below a threshold, the **Fusion Engine** will switch to the camera position which has higher accuracy. This will allow TerraFlex users to collect higher accuracy location data about their assets in GNSS degraded environments.

Mobility Data Visualization - A Tool for Detecting Suspicious Geospatial Patterns

R

Amy Sloane University of Canterbury

Supervisors: Vanessa Bastos, James Atlas

(Waimakariri 2, 9:30-9:45)

Geospatial researchers worldwide are very interested in the visualisation of mobility data. The potential of this data to reveal detailed patterns of human movement is significant, sparking extensive research in domains such as transportation. One domain remaining relatively unexplored is security and policing.

We present a practical solution for mobility data visualisation, potentially benefiting security and policing industries. We've developed a web-based application that detects suspicious or unusual activity, by highlighting residences with unusually high volumes of individuals passing through.

We followed a three-phased plan: research, development, and evaluation. We conducted systematic research on literature related to techniques and applications for visualising mobility data. This provided a solid foundation for the subsequent development phase. Our preliminary evaluation showed promising processing capabilities, efficiently handling 500 megabytes of data in approximately 30 seconds, and 1 gigabyte in roughly 1 minute and 20 seconds. Additionally, it identified suboptimal performance in 2 to 3 individual service methods.

These findings suggest that our solution has the potential to detect suspicious activity. We recognize the need for further improvements and are actively pursuing strategies such as data indexing and continued utilisation of computation libraries like Dask. In our final phase, we will conduct a thorough evaluation of performance using increasingly larger datasets. This will enable us to make informed predictions about the effectiveness of a fully functional, industry-deployed application.

Proof-of-location testing with a FOAM trust zone on UC campus

R

Jamie Thomas University of Canterbury, FOAM

Supervisor: Ben Adams

(Waimakariri 2, 9:45-10:00)

This project looks to set up a Trust Zone from FOAM on the University of Canterbury campus. This Trust Zone will provide a proof-of-location system similar to the Global Positioning System (GPS) in that it identifies nodes within a given area. This system uses Web3 and blockchain technology to make this identification decentralized and make it harder to spoof.

The Trust Zone is made up of four nodes, small boxes built on site and installed onto four rooftops across campus. Dynamic tests are run weekly to provide FOAM with data on how their technology is working in the new setting of UC. Dynamic tests involve moving around the campus with a FOAM mobile node connected to a device and using FOAM's web application, Hostal, record the node's location based on both GPS and the FOAM protocol.

The project also looks to implement a test bed web application to be used to record experiments using the FOAM protocol. These experiments can then be played back to answer questions about the FOAM technology, such as its capability of showing the location of a node compared to the Global Positioning System (GPS). The test bed will be built using React, NodeJS and Postgres/PostGIS. It is hosted on a virtual machine at the University of Canterbury to be used by testers. The test bed is not considered a commercial application to be used by future users of the FOAM Trust Zone, but instead by system testers. This application will record experiments alongside the Hostel web application from FOAM.

Mixed Reality Collaboration in Geoscience

Jack McCorkindale Seequent

Supervisor: Prof. Matthias Galster (Waimakariri 2, 10:00-10:15)

The use of mixed reality is rapidly becoming expected by key stakeholders in projects exceeding $100 million such as large infrastructure projects (e.g. dams, roads, rail) and projects on high stakes contaminated sites where geological understanding is important.

This project aims to develop an application that will allow greater collaboration in geoscience consultations through Mixed Reality using Microsoft HoloLens 2 devices. The application will allow the user to view a geological scene which represents the geographical location where construction occurs in 3D by scanning a QR code. The goal of the project is to enable geoscience consultants to collaborate with their clients in a way that improves understanding with minimal training required. A study performed on early access users provides insight into the effectiveness of the mixed reality application at meeting this goal.

Augmented Reality GIS Assets

I

Prableen Oberoi Trimble

Supervisor: Fabian Gilson

(Waimakariri 2, 10:15-10:30)

The project aims to address the challenge of transforming conventional 2D Geographic Information Systems (GIS) map views into an immersive 3D Augmented Reality (AR) experience. This transformation proves crucial in enhancing data acquisition and user interaction, notably for non-experts. The inception of the project saw the development of a Kotlin-based mobile app prototype, utilizing Google's ARCore Geospatial API and the device's GPS capabilities. This prototype furnished users with real-time location and position updates, allowing them to set markers on a 2D Google map, subsequently rendered as a 3D object in the AR view. Achieving horizontal accuracy up to 0.5 meters and vertical accuracy at 0.62 meters, the project then transitioned to the Xamarin platform, influenced by Trimble's extensive Xamarin usage which eases GIS data access. Moreover, anticipated integration with Trimble's Global Navigation Satellite Systems (GNSS) promises further enhancement in accuracy metrics. The finished Xamarin application now allows users to seamlessly track GIS assets in an AR display with real-time position and location updates on a 2D Google map. A key feature of the application is its ability to perform effectively even in locations with weak GPS signals, thanks to the ARCore Geospatial API which harnesses machine learning and Google Street View functionalities.

Session: Automated Machinery

(time: 11:00-12:15)

Chair: Chenyi Zhang

CNC Milling Chatter Detection Using a CNN AI

Lachlan Alsop Zealandia Systems, University of Canterbury

Supervisor: Fabian Gilson

(Waimakariri 2, 11:00-11:15)

Computer numerical controlled (CNC) milling is a subtractive manufacturing process that uses a spinning tool to remove material, it allows for complex and precise parts to be manufactured consistentaly and repeatadly. Self-excited chatter in CNC milling operations pose significant challenges, leading to tool and machine wear, increased power consumption, and decreased product quality. This project focuses on the detection of self-excited chatter, a critical step in mitigating its adverse effects. Chatter arises due to the forceful cutting interaction during milling, resulting in oscillations within the system. Traditionally, chatter detection relies on human ears and cognitive processing. However, this study introduces an approach employing a microphone and a convolutional neural network (CNN) to accurately predict chatter.

The solution is constrained by a processor orders of magnitude slower than the traditional commercial graphics processing units (GPU). This solution will be running on the OrangePi5, a system utilizing low power system architecture and some specialized CNN processing capabilities. In evaluation, the model demonstrates exceptional performance, achieving up to 95% accuracy on test data. As such this solution is optimized to run efficiently on the OrangePi5, achieving processing speeds ranging from 70 to 150 milliseconds using only the central processing unit. This capability allows for real-time analysis and further optimizations are to come.

Future work includes real-world testing, leveraging the specialized CNN processing capabilities of the OrangePi, and expanding the dataset to encompass various machine types. These endeavors aim to refine and generalize the chatter detection model, contributing to enhanced machining processes and product quality.

Insights Into Power Line Detection and Distance Estimation Using Stereo Cameras

R

Hugo Reeves Univeristy of Canterbury

Supervisor: Sam Schofield

(Waimakariri 2, 11:15-11:30)

Power line inspection utilising UAVs presents an opportunity to remove humans from a hazardous environment. Extensive previous research has shown success in power line detection and guidance techniques utilising hybrid LiDAR-Vision systems and event cameras. Due to the high cost of these systems, we sought to demonstrate a method for power line detection and depth estimation that leverages lower-cost fixed-frame-rate stereo cameras. Our proposed method attempts to detect power lines and estimate their depth; it does so using a stereo camera in combination with a Convolutional Neural Network (CNN) for power line edge detection and a line detection method previously only used with event cameras. Our CNN model was demonstrated to achieve comparable results to existing work. We discovered that the line detection method that had previously only been used with event cameras performs poorly with a low frame rate video stream. We build upon this conclusion by determining the qualities of this line detection method. We provide estimates of the error of depth estimation, relative to the frame rate, using this line detection method. To inform future research, we also provide an estimate of the equipment requirements neccesary for our proposed method to be applied successfully to a fixed-frame-rate camera.

Conductor Condition Recognition Using Autonomous Drones

I

Bede Skinner-Vennell Unison

Supervisors: Richard Green, Sam Scofield

(Waimakariri 2, 11:30-11:45)

Unison owns and operates the electricity network that distributes power to Hawke's Bay, Taupo, and Rotorua, which spans over 9,000km. Inspecting the condition of these conductors manually is a costly and time-consuming process. Unmanned aerial vehicles offer a promising alternative for inspecting power lines due to their advantages in safety, speed, and accuracy. This project aims to improve the efficiency of power line inspection by developing an autonomous drone system that can detect and follow power lines.

The proposed solution uses a LiDAR module to accurately detect the location of power lines in 3D space using unsupervised (rule based) classification and random sampling consensus (RANSAC). Our expected deliverable is an algorithm that outputs a vector that the drone should travel along in order to maintain the optimal position above the line for imaging purposes.

Smooth Drone Formation Flight

Euan Morgan, Frederik Markwell University of Canterbury

Supervisor: Andreas Willig (Waimakariri 2, 11:45-12:15)

Unmanned aerial vehicles, also known as drones have become feasible solutions to a number of real world issues. Several drones can be used together to enhance their usefulness. One major advantage of a formation of drones is that because of their separation, signals from the ground will reach them at different times. This allows for the swarm to triangulate the position of the signal much more accurately than a single drone. Possible applications of this technology include search and rescue, tracking of flying insects and firefighting.

Drones operating in formation require the ability to share position, movement and direction data to maintain the formation and avoid collisions. This must be done frequently and with minimal failures to ensure every drone has the most recent data possible. The created protocol addresses this requirement with a significant focus on maximising the frequency of communication. This is essential for project success as even slightly out-of-date data could cause a collision or failure within the formation. The protocol is developed in C++ and interfaces with the Boost and Robot Operating System libraries to provide communication between the protocol and other processes running on a particular drone.

A key difficulty encountered in previous work in this area is that flight tests are challenging and time consuming. It is difficult to configure the drones with the correct software, and to obtain real time information about the state of the drones while in the field. This is made more difficult by the lack of internet connectivity when away from base. The developed solution uses docker along with conventional linux tools such as scp to transfer software and configuration files. The backend was developed in Python, and the frontend in Javascript using ReactJS.

Session: Environmnent

(time: 13:15-14:30)

Chair: Trevor Nesbit

Scott Base Seal Census using Machine Learning

I

Jacques Terblanche Antarctica NZ

Supervisor: Richard Green

(Waimakariri 2, 13:15-13:30)

This project continues research that aids Antarctica New Zealand in monitoring the Weddell Seal population surrounding Scott Base.

A deep learning approach is used, employing a Convolutional Neural Network trained to detect seals in images captured from a camera overlooking Scott Base. This approach is able to count seals much faster and at more regular intervals than is possible manually, especially when there are regularly well over 200 seals in the area at one time. Analysing trends in data provided as a result of this project can help Antarctica NZ with identifying the colony's reactions to construction noise during the Scott Base reconstruction.

Mobility Data for Biosecurity

Karl Moore University of Canterbury Biosecurity Innovations

Supervisor: James Atlas

(Waimakariri 2, 13:15-13:30)

This project seeks to develop an application for monitoring individual travel patterns through bio-security risk zones and identifying potential risk hotspots. Researchers can draw specific regions of interest on a map and select custom time frames for analysis, with default polygons encompassing all New Zealand airports provided for convenience. Matching results are then displayed in the form of a heat map. Managing the substantial data volume posed challenges in storage and performance, leading to the development of a web-based application to eliminate individual user dataset management. To address computation efficiency, the project employs an R-tree for data organization, as well as sampling techniques to reduce server-to-client data payloads, ensuring responsive user experiences.

Using Mobile Location Data to Inform Climate-Resilient Transport Interventions

R

Jonathan Tomlinson Simon Kingham, University of Canterbury

Supervisors: James Atlas, Vanessa Bastos

(Waimakariri 2, 13:30-13:45)

New Zealand has a need to reduce its greenhouse gas emissions in order to meet its international obligations as well as targets in domestic legislation. As road transportation makes up 15% of New Zealand's total greenhouse gas emissions, significant changes will need to be made to the country's transport infrastructure to achieve the required cuts in emissions. Shifting away from private motor vehicles and towards public and active transportation is a key component of these changes.

Current methods to evaluate the effectiveness of transport interventions rely on the manual counting of vehicles, cyclists, and pedestrians. Such methods can be slow, costly, and inaccurate. This project leverages large amounts of mobile phone location data to evaluate existing transport interventions and improve future ones. Mobile phone location data can provide a clear picture of transport habits over longer time spans and in more locations than data obtained via traditional methods, such as traffic on nearby roads.

The proposed solution allows transport researchers to view traffic volumes at the level of street segments, and compare volume between different time periods. For example, researchers can view traffic volumes from before and after a particular intervention to assess its overall effectiveness. The insights gained from this data can then be used to make more effective transport changes in the future.

Port Identifier for Pests (PIP)

Lane Edwards-Brown AgResearch

Supervisor: James Atlas

(Waimakariri 2, 13:45-14:00)

The Port Identifier for Pests (PIP) is a cross-platform mobile application intended to support individuals involves in port operation in protecting our biosecurity.

Protecting our natural environment from invasive pests and diseases is a key part of border security, with shipping contianers serving as a vector for these threats to enter New Zealand, PIP seeks to assist in reporting these threats in order to aid in management on the local level (helping worker know how to react to an unknown species) and the national level (helping provide information on possible pest encounters).

By incorporating modern app design with expertise knowledge on pests and related species, PIP allows port operators to quickly incorporate pest identification and reporting into their typical work-flow - encouraging them to contribute to biosecurity. This is achieved through quick and simple actions and the addition of a database of insects that is pre-sorted based on the location of the user and the incoming shipment.

Efficient Annotation of Sparse Vocalisations for Bioacoustic Monitoring

R

Shaylin Simadari University of Canterbury

Supervisor: Andrew Bainbridge-Smith

(Waimakariri 2, 14:00-14:15)

Bioaccoustic monitoring involves recording audio from ecosystems and estimating species density based on the frequency of animal vocalisations. The aim of this project is continue developing a tool for annotating vocalisations and training ML models for automatic detection. Domain experts that use annotation tools, may not be experienced in using nongraphical user interfaces, as such creating a user friendly graphical interface may be integral to providing a means for quick annotation."

Session: Varia

(time: 15:00-16:15)

Chair: Marina Filipovic

Sentiment analysis for licensing applications

R

Oliver Johnson CSSE

Supervisor: Ben Adams

(Waimakariri 2, 15:00-15:15)

This project concerns itself with alcohol licensing decisions in Christchurch. Members of the public may object to an alcohol licence, but in approximately 80 percent of cases the licence is granted despite the objections. Many objectors (and their communities) feel disheartened from this and are left wondering how they can make an effective objection.

This project aims to investigate if machine learning techniques are able to illuminate why most licences are granted and possibly even provide guidance on making more effective objections.

The methods used include naive bayesian analysis and LDA topic modelling. Both of these methods respond to patterns within their data, which is well-suited to the task of illuminating such patterns. In particular, topic modelling was used extensively (and with many variations) to analyse the licensing decisions.

Digital Twin Builder

I

Max Bastida Eagle Technology

Supervisor: Ben Adams

(Waimakariri 2, 15:15-15:30)

An application was developed to generate digital twins of New Zealand cities for use in urban planning and addressing heat islands in urban areas. In collaboration with Eagle Technology, the ArcGIS platform was used to create an application that could generate three dimensional models of New Zealand cities from location data for trees and buildings. These models also display data about the location and intensity of urban heat islands. Urban heat islands occur because heat is generated and trapped by materials and morphology of urban areas. Urban planners can edit the trees and buildings in the model which will update to show the effect of their changes on urban heat islands. This will help urban planners visualise areas and what adjustments can be made, and how this will impact the heat that is trapped within the city.

Bloodstain Pattern Classification using Deep Learning

I

Aidan Campbell Institute of Environmental Science and Research

Supervisors: Andrew Bainbridge-Smith, Richard Green, Rosalyn Rough (Waimakariri 2, 15:30-15:45)

Bloodstain pattern analysis is the practice of investigating bloodstain patterns from crime scenes and accidents to assist in evidence gathering. This practice is subject to indiviual bias, differences in experience, and human error, therefore a method is needed to quantitatively analyse a bloodstain pattern.

This project aims to assist forensic scientists by providing a method for bloodstain pattern analysis using deep learning. Building on the work of previous students, the objective of this project was to explore how effective neural networks which are trained with the image data of bloodstains are at identifying three patterns.

A network was trained using information such as the circularity, angle and intensity of each stain within the overall pattern. A second solution was investigated where the coordinates of the boundary pixels of bloodstains were input into a point cloud neural network called PointNet++. This network achieved high accuracy that was similar to or better than previous work.

Laboratory automation at ESR

I

Moses Wescombe The Institute of Environmental Science and Research (ESR)

Supervisor: Fabian Gilson

(Waimakariri 2, 15:45-16:00)

This project proposes a design and method that implements an automation system for gamma spectroscopy. The system is designed to increase the throughput of the gamma spectroscopy lab at The Institute of Environmental Science and Research (ESR).

Currently, the lab is operated manually during limited hours. The proposed design uses a robotic arm combied with computer vision to move the samples in and out of the detectors automatically, allowing for 24-hour operation. The operator will manage the processing of samples through a web app that communicates with the robotic arm via a server connection.

Thanks to our sponsors

Our sponsors make such events possible. We truly appreciate their support. If you would like to be a sponsor for Computer Science and Software Engineering events, please contact **csse@canterbury.ac.nz**.

Photos and videos

Photos and/or video will be taken at this event. The official photos/videos will only be used for promoting the university and its courses. Please let the event co-ordinator know if you do not want your photo taken so that the photographers will avoid to take pictures of you and blur your face in the final product, if necessary.

Credits

This booklet has been prepared by Etienne Borde and has been adapted from the work of Fabian Gilson, Maxime Lucas, Pau Clusella and Thomas Kreuz.
The cover of this booklet has been designed by Etienne Borde, using as background a picture designed by Harryarts / Freepik.

Source files available at **`https://gitlab.com/fabgilson/conference-booklet`** (under GPL3 license).

UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Te Tari Pūhanga Pūmanawa Rorohiko
Computer Science & Software Engineering


csse@canterbury.ac.nz


uc.csse


@uccsse


uccsse