

Software Engineering Final Year Projects Showcase

Rydges Latimer Hotel
Ōtautahi Christchurch

2022



Computer Science & Software Engineering

SOFTWARE ENGINEERING FINAL YEAR PROJECTS

Showcase

Rydges on Latimer Hotel, Ōtautahi Christchurch (Aotearoa New Zealand)

14th October 2022



Foreword	7
Programme	9
TRACK: Charles Babbage	15
MACHINE LEARNING <i>Chair: Ben Adams (9:30-10:30)</i>	15
Application of Convolutional Neural Networks on Bloodstain Pattern Analysis	
<i>Swapnil Bhagat</i>	15
Barracuda - A Stack Based Programming Language Compiler For GPUs	
<i>Connor Macdonald</i>	16
Automated Extraction of Jumping Spider Behaviour from Experiment Footage	
<i>Jack Patterson</i>	17
SICLOPS - An AI for the Creation of Generalised Datasets	
<i>Dan Chapman</i>	18
ENVIRONMENT <i>Chair: James Atlas (11:00-12:15)</i>	19
Estimating annual greenhouse gas emissions from different working scenarios	
<i>Ronan Avery</i>	19
Virtual reality interaction for apple thinning training tool	
<i>Rutger van Kruiningen</i>	20
Vineyard Visualisation	
<i>Niels van Antwerpen</i>	21
How to count 332 seals in two seconds	
<i>Fletcher Dick</i>	22
3D Model For Digital-Twin of a Plant	
<i>Inga Tokarenko</i>	23
DATA VISUALISATION <i>Chair: Andreas Willig (13:15-14:30)</i>	24
Sequent Siteviz: Experiencing Geology in-situ	
<i>Rio Ogino</i>	24

CameraX - Next Generation of Field Capture	
<i>Nitish Singh</i>	25
Environmental Data Visualisation Tool	
<i>Darcy Maddren</i>	26
Digital twins with Visual Simultaneous Localization & Mapping	
<i>Henry Barrett, Joseph Kelly</i>	27
UNMANNED VEHICLES <i>Chair: Fabian Gilson (15:00-16:00)</i>	28
Automated UAV Pruning of Pinus Radiata Catkins	
<i>Eric Song</i>	28
Underwater blue light (BiFi) comms for underwater drones	
<i>James Harris</i>	29
UAV Formations	
<i>Caleb Sim, Niko Tainui</i>	30
TRACK: Ada Lovelace	31
SOFTWARE ENGINEERING <i>Chair: Marina Filipovic (9:30-10:30)</i>	31
Towards Understanding Communication in Software Engineering Team Meetings	
<i>Joshua Egan</i>	31
Software Quality Tools in the Wild	
<i>Michael Morgoun</i>	32
Gamification in Educational and Industrial Contexts: Play and Learn (Query Competition)	
<i>Edward Langlands</i>	33
Gamified Scrum Retrospectives to Enhance Motivation and Engagment in Learning	
<i>Seth Kingsbury</i>	34
EDUCATION <i>Chair: Miguel Morales-Trujillo (11:00-12:15)</i>	35
Implement automated deployment, monitoring and logging to a Docker swarm	
<i>Samuel Burtenshaw</i>	35
Event Registrations for DTHM4Kaiako Web Application	
<i>Hayley Krippner</i>	36
Skill Tracking and Recommendations for CodeWOF	
<i>Griffin Baxter</i>	37

Recording Design Decisions in Slack	
<i>Zachary Kaye</i>	38
Question tracking system for Moodle	
<i>Gongzai Li</i>	39
E-HEALTH <i>Chair: Clémentine Gritti (13:15-14:30)</i>	40
Measuring bone health features in cortical and trabecular bone with minimal user intervention	
<i>Oliver Cranshaw, Michael Woodard</i>	40
Retinal image quality assurance using deep learning technique	
<i>Nathan Huynh</i>	41
Tumour Region Segmentation in Whole Slide Images	
<i>Mitchell Veale</i>	42
3d Reconstructions of Tumour Resections Using Whole Slide Images	
<i>Mitchell Freeman</i>	43
CRYPTOGRAPHY <i>Chair: James Atlas (15:00-16:00)</i>	44
Differential Privacy Tools for Smart Contracts	
<i>Kye Oldham</i>	44
Post-Quantum DTLS and MQTT protocols	
<i>Callum McLoughlin</i>	45
Post-Quantum DTLS and MQTT protocols	
<i>Raven Townsend</i>	46
Ternary Merkle Trees for Use in the Internet of Things	
<i>Finn Bright</i>	47
THANKS TO OUR SPONSORS	49

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 their preparedness 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, SENG402 course coordinator

On behalf of the COMPUTER SCIENCE AND SOFTWARE ENGINEERING DEPARTMENT

SUMMARISED SCHEDULE

	Savoy East	Clarendon
8:45-9:15	Coffee & Tea	
9:15-9:30	Welcome Words (Savoy East)	
9:30-10:30	Machine Learning chair: Ben Adams	Software Engineering chair: Marina Filipovic
10:30-11:00	Morning Tea	
11:00-12:15	Environment chair: James Atlas	Education chair: Miguel Morales-Trujillo
12:15-13:15	Lunch	
13:15-14:30	Data Visualisation chair: Andreas Willig	E-Health chair: Clémentine Gritti
14:30-15:00	Afternoon tea	
15:00-16:00	Unmanned vehicles chair: Fabian Gilson	Cryptography chair: James Atlas
16:00-16:15	Closing Words (Savoy East)	
16:15-17:30	Social Event	

PRESENTATION COLOURED BADGES



Industry project




Research project



Personal project

8:45-9:15		Coffee & Tea	
9:15-9:30		Welcome Words (Savoy East)	
9:30-9:45		Swapnil Bhagat <i>Institute of Environmental Science and Research</i>	Application of Convolutional Neural Networks on Bloodstain Pattern Analysis
9:45-10:00		Connor Macdonald <i>University of Canterbury, Te Whatu Ora Waitaha Canterbury</i>	Barracuda - A Stack Based Programming Language Compiler For GPUs
10:00-10:15		Jack Patterson <i>University of Canterbury</i>	Automated Extraction of Jumping Spider Behaviour from Experiment Footage
10:15-10:30		Dan Chapman <i>Computer Concepts Limited</i>	SICLOPS - An AI for the Creation of Generalised Datasets
10:30-11:00		Morning Tea	
11:00-11:15		Ronan Avery <i>Enable Networks Limited</i>	Estimating annual greenhouse gas emissions from different working scenarios
11:15-11:30		Rutger van Kruiningen <i>University of Canterbury, Maaratech</i>	Virtual reality interaction for apple thinning training tool
11:30-11:45		Niels van Antwerpen <i>University of Canterbury, Maaratech</i>	Vineyard Visualisation
11:45-12:00		Fletcher Dick <i>Antarctica NZ</i>	How to count 332 seals in two seconds
12:00-12:15		Inga Tokarenko	3D Model For Digital-Twin of a Plant
12:15-13:15		Lunch	
13:15-13:30		Rio Ogino <i>Sequent</i>	Sequent Siteviz: Experiencing Geology in-situ
13:30-13:45		Nitish Singh <i>Trimble</i>	CameraX - Next Generation of Field Capture
13:45-14:00		Darcy Maddren <i>University of Canterbury</i>	Environmental Data Visualisation Tool
14:00-14:30		Henry Barrett, Joseph Kelly <i>Beca</i>	Digital twins with Visual Simultaneous Localization & Mapping
14:30-15:00		Afternoon tea	
15:00-15:15		Eric Song <i>University of Canterbury</i>	Automated UAV Pruning of Pinus Radiata Catkins
15:15-15:30		James Harris <i>Cawthron Institute</i>	Underwater blue light (BiFi) comms for underwater drones
15:30-16:00		Caleb Sim, Niko Tainui <i>University of Canterbury</i>	UAV Formations
16:00-16:15		Closing Words (Savoy East)	
16:15-17:30		Social Event	

8:45-9:15		Coffee & Tea	
9:15-9:30		Welcome Words (Savoy East)	
9:30-9:45		Joshua Egan <i>University of Canterbury</i>	Towards Understanding Communication in Software Engineering Team Meetings
9:45-10:00		Michael Morgoun <i>University of Canterbury</i>	Software Quality Tools in the Wild
10:00-10:15		Edward Langlands <i>University of Canterbury</i>	Gamification in Educational and Industrial Contexts: Play and Learn (Query Competition)
10:15-10:30		Seth Kingsbury <i>University of Canterbury</i>	Gamified Scrum Retrospectives to Enhance Motivation and Engagement in Learning
10:30-11:00		Morning Tea	
11:00-11:15		Samuel Burtenshaw <i>UC Computer Science Education Research Group</i>	Implement automated deployment, monitoring and logging to a Docker swarm
11:15-11:30		Hayley Krippner <i>UC Computer Science Education Research Group</i>	Event Registrations for DTHM4Kaiako Web Application
11:30-11:45		Griffin Baxter <i>UC Computer Science Education Research Group</i>	Skill Tracking and Recommendations for CodeWOF
11:45-12:00		Zachary Kaye <i>University of Canterbury</i>	Recording Design Decisions in Slack
12:00-12:15		Gongzai Li <i>University of Canterbury</i>	Question tracking system for Moodle
12:15-13:15		Lunch	
13:15-13:45		Oliver Cranshaw, Michael Woodard <i>MARS Bioimaging</i>	Measuring bone health features in cortical and trabecular bone with minimal user intervention
13:45-13:00		Nathan Huynh <i>University of Canterbury, Te Whatu Ora Waitaha Canterbury</i>	Retinal image quality assurance using deep learning technique
14:00-14:15		Mitchell Veale <i>University of Canterbury</i>	Tumour Region Segmentation in Whole Slide Images
14:15-14:30		Mitchell Freeman <i>University of Canterbury</i>	3d Reconstructions of Tumour Resections Using Whole Slide Images
14:30-15:00		Afternoon tea	
15:00-15:15		Kye Oldham <i>University of Canterbury</i>	Differential Privacy Tools for Smart Contracts
15:15-15:30		Callum McLoughlin <i>University of Canterbury</i>	Post-Quantum DTLS and MQTT protocols
15:30-15:45		Raven Townsend <i>University of Canterbury</i>	Post-Quantum DTLS and MQTT protocols

15:45-16:00		Finn Bright <i>University of Canterbury</i>	Ternary Merkle Trees for Use in the Internet of Things
16:00-16:15		Closing Words (Savoy East)	
16:15-17:30		Social Event	

SESSION: MACHINE LEARNING

(time: 9:30-10:30)

Chair: Ben Adams

Application of Convolutional Neural Networks on Bloodstain Pattern Analysis

**Swapnil Bhagat** Institute of Environmental Science and Research

Supervisor: Rosalyn Rough

(Savoy East, 9:30-9:45)

Bloodstain Pattern Analysis (BPA) is the forensic discipline of analysing blood spatter patterns at crimes scenes and other blood letting events.

The analysis is often done manually and qualitatively, introducing human error and subjectivity. With error rates of 13-23%, analysts may benefit from

more quantitative and objective approaches such as in machine learning. Building on previous works that research various machine learning models, this project explored more specifically the applicability of Convolutional Neural Networks

(CNNs) on the classification of *expired*, *cast-off*, and *impact* bloodstain patterns. In particular, the project aimed to answer: How effective are CNNs in classifying complete and partial bloodstain spatter patterns.

To achieve this, the most important and first step was the creation of a repository of datasets to be used for training and testing CNN models. As the raw images were large and varied in size, it was computationally unfeasible to train them directly. This raw dataset was manipulated (e.g., dilated, resized, cropped, etc.) using preprocessing techniques to create the aforementioned repository. Training datasets were then fed into pretrained models - VGG16 and ResNet50, to produce re-usable and testable models. Results indicate that indeed CNNs are a feasible solution for BPA analysts to classify complete or near-complete bloodstain patterns (>95% accuracies), but struggle with partial patterns (<90% accuracy).

Barracuda - A Stack Based Programming Language Compiler For GPUs



Connor Macdonald University of Canterbury, Te Whatu Ora Waitaha Canterbury

Supervisor: Phillip Duncan-Gelder

(Savoy East, 9:45-10:00)

Graphical Processing Units (GPUs) were traditionally developed to manage the highly distributed workload of rendering graphics to high resolution modern monitors. Because of their wide availability and parallel computation design, researchers over the last decade have been investigating their ability for alternative use cases. This has largely been a successful endeavor and has accelerated new technologies such as machine learning and large scale simulation models. Demand for GPU software has flourished in order to take advantage of the improvements parallel computing offers.

Currently to leverage the GPU, researchers must either write in a low level language such as C, a task which is difficult for those who are not native programmers, or use a domain specific API that may not be versatile enough. This impedes the speed at which research is conducted and can create an unnecessary barrier to entry, especially for work in unrelated domains.

Barracuda is a high level language designed to be run and executed on the GPU. Being a high level language it is able to offer a simple interface and model for writing and running code on a GPU. The collection of tools in providing this language aims to provide researchers with minimal barriers when conducting research, while democratizing programming on the GPU. Additionally the toolkit has been written to allow simulation framework developers an easy user scripting language to hook into their models.

A key component of the Barracuda language is the development of a compiler. The compiler is responsible for generating the lower level instructions provided to the Barracuda runtime from a Barracuda source file. The achieved purpose of this project was to develop this compiler which was capable of modern language features such as variables, branching conditionals, loops, and functions.

Automated Extraction of Jumping Spider Behaviour from Experiment Footage



Jack Patterson University of Canterbury

Supervisors: Dr James Atlas, Prof Ximena Nelson

(Savoy East, 10:00-10:15)

Biologists at the School of Biological Sciences at the University of Canterbury are currently studying the behaviour of jumping spiders to develop more accurate models of their cognitive and visual responses. This research involves the use of choice test experiments, which are recorded and manually examined by an observer. This is a time-consuming process as it involves watching hundreds of experiment videos and noting down all instances of the behavioural features the researchers wish to study. This project accelerates the footage examination process by automatically performing subject tracking and pose estimation on experiment footage and using that data to extract the behavioural features normally supplied by a human observer.

To accomplish this, a three-stage solution was designed. The first stage was developed by choosing the most accurate approach out of three alternative subject tracking methods tested on our dataset. This stage uses an object detection network to detect the subject in each frame, along with a second network to track the same identity across frames. Similarly, for the second stage, three possible alternatives were examined to determine the best technique to use for animal pose tracking for the relevant jumping spider species. The second stage uses a segmentation network to produce the locations of key points on the spider's body for each frame. By combining and analysing the outputs of these two stages, the third stage programmatically extracts the behavioural features present in each video, which are then displayed to the user.

SICLOPS - An AI for the Creation of Generalised Datasets



Dan Chapman Computer Concepts Limited

Supervisors: William Fulton, Josh Hobbs

(Savoy East, 10:15-10:30)

A problem with establishing new computer vision projects is the need for a suitably labelled or segmented dataset to train the model on – creating such a dataset often requires a large number of person hours, with the correct segmentation of potentially thousands of images taking days to be approved. SICLOPS (Sets of Images Created and Labelled through Object Perception and Segmentation) is a computer vision product that aims to help machine learning and computer vision developers establish a training dataset quickly.

SICLOPS attempts to automate the dataset creation process – a small number of images are hand-labelled, and are then fed into the model as training data. The model then attempts to label new data, and presents these to the user, who marks these attempts on their suitability. Suitable images are then fed back into the model as additional training data, so that the model can learn from its successful attempts at predictions.

This process aims to result in a fully-trained model, that can successfully label new images passed to it. This would save developers the large time-sink of creating the training dataset, as only a small number of images would need to be created to establish the model, and the amount of human interaction to generate a large dataset would be replaced by the computation of the models. Initial tests show the model has been able to achieve an accuracy of 91.07% labelling images and 86.17% accuracy segmenting images, with only 50 initial images and 4 training runs.

SESSION: ENVIRONMENT

(time: 11:00-12:15)

Chair: James Atlas

Estimating annual greenhouse gas emissions from different working scenarios



Ronan Avery Enable Networks Limited

Supervisor: Melissa Keys

(Savoy East, 11:00-11:15)

With a recent increase in internet traffic globally and many countries joining the Paris Agreement, a need for sustainable practices has also increased. Organisations and businesses are striving to create sustainable working environments.

This project is sponsored by Enable Networks and aims to introduce a model to estimate the greenhouse gas emissions created when combining working from home and working from the office. This model will be actualized through a user-friendly online tool used to calculate and compare emissions from different working scenarios.

By providing a tool that produces quantifiable values, awareness of the sustainability of different working environments will be increased, so that people can consciously help lower our net carbon footprint.

Virtual reality interaction for apple thinning training tool



Rutger van Kruiningen University of Canterbury, Maarattech

Supervisor: Prof Richard Green

(Savoy East, 11:15-11:30)

This project aims to determine preferred mechanics for interacting with small objects, namely apple fruitlets in a virtual reality environment using hand tracking.

An apple thinning training game is being developed in conjunction with this project which requires these apple interaction mechanics to provide an immersive and uninterrupted experience for the user.

For this, multiple interaction candidates were developed and then tested by various users online in hopes to determine the most immersive apple thinning experience after taking into account data from these users' performance and accuracy.

Initially, research was conducted into existing solutions for interacting with objects in the virtual world including attaching mechanisms to the user's fingers to simulate resistance, but using the built-in hand tracking on the Meta Quest was decided to be used due to its lightweight nature and easy addition to the project codebase.

Vineyard Visualisation



Niels van Antwerpen University of Canterbury, Maaratech

Supervisor: Prof Richard Green

(Savoy East, 11:30-11:45)

There is currently no easy method of retrieving historical data about a particular plant when in a vineyard. This project proposes a simple workflow for solving this data retrieval problem. To find a particular plant while in a vineyard a mobile application was created. This application is used to collect and upload images of the target plant. GPS data is used to increase accuracy of comparison and decrease the size of plants that are compared.

The uploaded images are then segmented to isolate the plant and compared with previous images. Comparing the images is done using two computer vision algorithms, to detect and compare features in the images the Oriented FAST and Rotated BRIEF (ORB) algorithm was used, and to detect structural similarity in the images the algorithm Structural Similarity Index Metric (SSIM) was used. These are used to identify the correct plant from the database. The application can show the historical data from the now identified plant, this has been tested with data collected with Maaratech's scanning platform.

How to count 332 seals in two seconds



Fletcher Dick Antarctica NZ

Supervisor: Dr Oliver Batchelor

(Savoy East, 11:45-12:00)

This project aims to conduct research that allows Antarctica New Zealand to monitor the Weddel Seal population surrounding Scott Base. Frequent population counts will allow Antarctica NZ to identify colony disturbance caused by the upcoming Scott Base reconstruction. Surrounding Seal counts can exceed 300, making frequent manual counting laborious and inaccurate due to mental fatigue.

Such a task lends itself well to a Convolutional Neural Network (CNN). A CNN is a deep learning algorithm that can learn to analyse imagery. The proposed solution uses a CNN for its object detection capabilities. An existing trained CNN was selected for its similarity in purpose and then re-trained using images captured from a high-resolution camera overlooking Scott Base. This Transfer Learning technique resulted in an accurate and robust solution to Seal monitoring.

3D Model For Digital-Twin of a Plant



Inga Tokarenko

Supervisor: Prof Richard Green

(Savoy East, 12:00-12:15)

The project aims to assist a solution to the current over usage of water in our agricultural industries.

The project focuses on creating a 3D model for a digital-twin of a tomato plant and expected growth of the plant over a set period.

The solution provides a utilized image processing technique to better the digital-twin, computer vision and 3D reconstruction techniques that supported better outcomes in providing less noise, higher density of the model and height measurements with accuracy within five or less centimetres.

SESSION: DATA VISUALISATION

(time: 13:15-14:30)

Chair: Andreas Willig

Seequent Siteviz: Experiencing Geology in-situ



Rio Oginio Seequent

Supervisor: Tim Schurr

(Savoy East, 13:15-13:30)

Seequent is a leading software company in the geosciences space which provides geoscientists with powerful desktop tools to create 3D geologic models. However, geoscientists have had no way of bringing these models to the field other than static, 2D cross-sections - vertical slices through the earth which expose the geology hidden underground.

For the past several years, Seequent has been working towards a solution to this problem in the form of an iOS app, Seequent Siteviz. With the app nearing a releasable state, this project has focused on developing features in collaboration with geoscientists, using the insights that only they can bring to drive the development process.

Siteviz allows users to create an interactive, vertical cross-section of the model under their feet, using built-in or external sensors to determine their current position and orientation. With this, users can identify features such as drill holes, transition depths and hazardous minerals such as asbestos in relation to themselves. Compared to paper printouts or PDFs, Siteviz will give users a more accurate and relevant view of the scene, making models more approachable and understandable to geoscientists, as well as other on-site stakeholders who may have limited domain knowledge.

CameraX - Next Generation of Field Capture



Nitish Singh Trimble

Supervisors: Guillaume Clin, Stuart Ralston

(Savoy East, 13:30-13:45)

CameraX is a R&D project proposed by an American industrial technology company - Trimble. Currently, Trimble's customers use instruments that are equipped with high-end camera, IMU (Inertial Measurement Unit) sensor and GPS antenna to capture field images and the required metadata for photogrammetry. These high precision and expensive instruments require deep knowledge to operate. This results in more time, effort and money due to the setup, operation and maintenance by the highly trained professionals.

This project aims to minimise the challenges faced by Trimble's customers by automating the image capturing and metadata gathering process for photogrammetry. As part of this project, an android application has been developed. This application has a user-friendly user interface which allows the user to specify a frequency and automatically capture high quality images based on that frequency. These images also include device location and orientation information which is embedded into the image files using Adobe's file labeling technology - Extensible Metadata Platform (XMP). The application also allows the users to customise the image quality by controlling the various camera parameters like focus, shutter speed, ISO and metering with a user-friendly user interface.

This application was tested for usability by getting four Trimble employees to interact with the application and observing their behaviour and reactions to it. The results of the usability study suggested that the application met the goal of automating the image capture and metadata gathering process as well as controlling the image quality. However, the results also suggested that the user interface could be improved for better user experience.

Environmental Data Visualisation Tool



Darcy Maddren University of Canterbury

Supervisors: A/Prof Ben Adams, Marwan Katurji, Jaiwei Zhang

(Savoy East, 13:45-14:00)

This project aims to provide researchers and others with tools to explore data from the New Zealand Modelling Consortium Open Environmental Digital Library (Envlib). With Envlib's mission to supply free weather and climate model data, the problem space expanded to assisting users who may not have the tools or expert knowledge to visualise or otherwise explore the data.

Therefore, a solution to visualise Envlib datasets in a public web application was proposed. The solution uses a full stack approach from preprocessing to client-side visualisation. When requests are sent to the Envlib API, data from the Tethys python package is first preprocessed into an optimised format. On return this is sent back to a React application which renders the results as a set of visualisation layers on an interactive web map of New Zealand. Currently these visualisations show dynamic wind vectors and heat maps.

Digital twins with Visual Simultaneous Localization & Mapping



Henry Barrett, Joseph Kelly Beca

Supervisor: Sam Houghton

(Savoy East, 14:00-14:30)

FACILITYtwin[®], a digital twin platform developed by Beca, helps owners and managers of manufacturing and utilities facilities get easy access to the data they need to make better decisions about their assets. An important part of this is the contextualization of assets using 360-degree photo virtual tours. The creation of 360 tours is limited because it requires a manual error-prone process to provide context for understanding how the images relate to one another, and the virtual tours are "flat" in that they use 2-dimensional 360 degree photos.

The goal of the first subproject is to automatically generate a photo tour of a facility from a set of 360 degree photos. Users can transition between adjacent photos in a 3D environment, and locations of each photo can be seen on a map of the facility. The proposed solution implements Visual Simultaneous Localization and Mapping (VSLAM) through a modified Stella-vslam system. Stella-vslam is a system that can perform VSLAM with equirectangular images as input. The system has successfully identified the locations of photos in an environment using video input, with post-processing to identify the key photos (video frames) used in the photo tour. Further research and modification is required for the system to consistently identify locations of photos that are separated by a distance larger than approximately 1.4m (the current limit).

The goal of the second subproject is to generate a 3D model of a scene from multiple 360 degree photos. 360 degree photos are decomposed into multiple perspective images. The perspective images are then fed through a pre-processing step to determine their relative positions to one another. A technique called neural radiance fields is used to train a neural network on the images to generate a 3D model of the scene. The model can be cropped and viewed with six degrees of freedom to provide unique top-down viewpoints, greater context, and immersion.

Automated UAV Pruning of Pinus Radiata Catkins**Eric Song** University of Canterbury

Supervisors: Prof Richard Green, Sam Schofield

(Savoy East, 15:00-15:15)

Pinus Radiata trees form pollen-producing catkins which can be harvested for pharmaceutical uses. As P. Radiata is also the most common tree species grown in New Zealand forestry, there is potential for a lucrative catkin-harvesting industry. However, current manual harvesting methods are not viable in most forests, as P. Radiata are grown to heights of over 20 meters.

Unmanned Aerial Vehicles, or drones, could be well suited to harvesting catkins at these heights. We present a method to reliably detect and track catkins, providing real time guidance for an autonomous harvesting drone. The proposed method utilises a custom-trained YOLOv5-nano deep learning model to detect each catkin. Image segmentation techniques are used to extract the 3D location and shape of the catkin. A RANSAC-assisted variant of the Kanade-Lucas-Tomasi tracker provides real-time tracking of catkins during flight. The method has been integrated onto a harvesting drone and successfully guided the drone to harvest catkins from a P. Radiata tree.

Underwater blue light (BiFi) comms for underwater drones



James Harris Cawthron Institute

Supervisor: Prof Andreas Willig

(Savoy East, 15:15-15:30)

Autonomous underwater drones for use in aquacultural settings are being developed at the University of Canterbury. A requirement for the drones is to send video data back to an operator on the surface. A key constraint is that the communication must be done in a wireless fashion as the drones will be operating in environments such as mussel farms which have underwater ropes and could become tangled if tethered.

An existing system uses blue LEDs to transmit/receive data underwater, referred to as BiFi. This system makes use of a packetization protocol that splits the data into packets, which among other fields contains a checksum so that the integrity of the data can be verified on the receiver end.

This project aims to make the delivery of the video data more robust by extending the packetization protocol to implement medium access control and link layer protocols to support standard error discovery and repair.

The proposed solution uses a combination of the Selective Repeat ARQ protocol and ALOHA protocol to increase robustness to errors. Selective Repeat ARQ introduces re-transmissions for packets that were damaged or lost in transmission increasing reliability as data is less likely to be truly lost. The ALOHA protocol introduces a short random backoff after failed transmissions which decreases the chances of a subsequent transmission also failing.

Turbidity is the measure of relative clarity of a liquid and causes scattering of light that is shined through the water. This scattering is the main cause of errors when transmitting data through water using blue light. The protocol has been tested through a set of experiments to determine its reliability and performance at a range of turbidity levels and distances.

UAV Formations



Caleb Sim, Niko Tainui University of Canterbury

Supervisors: Prof Andreas Willig, Sam Pell, Dr Graeme Woodward

(Savoy East, 15:30-16:00)

Autonomous drone formations allow the ability to gather information in terrain usually hard to reach for humans. Whether this is a search and rescue operation in the mountains or tracking insects to greater learn about their species. These autonomous formations have many requirements to achieve their purpose. Some of these requirements include stable network communication, clear internal architecture and a robust safety framework.

This undertaking required a custom network layer protocol that allows communication between drones and a ground control station with specific information. The size of this project in terms of complexity and developer size requires a clear architecture. This clear architecture will facilitate the expansion of the features of the formation by future developers. Drones are capable of high-speed movement which is a risk to not only other drones in the formation but also their surroundings. The solution to this issue was a reliable safety framework to reduce the risk of such a complex operation.

The safety mechanisms implemented in the formation include collision avoidance and a geofence. Since these drones are operating in formation it is of the utmost importance that there is a clear line of communication especially if one falls into a state of error. The drones consist of a PX4 flight controller which receives signals from a raspberry pi that runs our code. The drones communicate over a wireless network using a custom network layer protocol.

SESSION: SOFTWARE ENGINEERING

(time: 9:30-10:30)

*Chair: Marina Filipovic***Towards Understanding Communication in Software Engineering Team Meetings****Joshua Egan** University of Canterbury

Supervisors: Prof Tanja Mitrovic, A/Prof Matthias Galster

(Clarendon, 9:30-9:45)

Software engineers, as with many other professions, need to be able to communicate effectively with members of their team. They will also spend large amount of time in meetings, which is where critical decisions will be made and the organisation of the team and their current project will be arranged. Hence, the ability to communicate in team meetings is a vital skill for a software engineer, and to improve upon this skill first requires a deeper analysis of what happens in meetings. This project aims to better understand the behaviours exhibited by Software Engineering students in such meetings.

To better understand communication in meetings, 10 recordings of SENG202 teams who were asked to demonstrate a normal meeting for their course project. These were analysed using Epistemic Network Analysis (ENA), a process that combines Qualitative and Quantitative research to create a network of behaviours to compare. This process means that not only are behaviours exhibited tallied, but also displayed as connections to other behaviours that often occurred alongside them, giving a much better visual representation of each team's activities and focus.

This project aimed to analyse the behaviours exhibited based on the productivity, engagement, and length of each team meeting to find common points that should be strived for or avoided. The use of ENA enabled the comparison of productive and unproductive teams in terms of crucial elements. These differences were not immediately apparently without the use of ENA.

Software Quality Tools in the Wild



Michael Morgoun University of Canterbury

Supervisor: Dr Miguel Morales-Trujillo

(Clarendon, 9:45-10:00)

Software quality is a crucial factor in the development of software for the industry. It can impact not only the functional use of the software but also the satisfaction of its users and its longevity. While critical, in most cases, software quality is not defined, quality attributes are not specified, and the quality tools and practices used during development remain undetermined.

To fill this gap, this project aims to determine how software practitioners define software quality and what tools/practices they use to uphold their software quality.

This project consisted of two phases, a preliminary research phase to find appropriate literature on the usage of software quality tools in the industry, and a data collection phase.

The method used in the data collection phase was to conduct semi-structured interviews with experienced and currently employed software practitioners alongside distributing a survey to all software practitioners with any experience in the software industry. The interviews contained a predetermined list of questions but through the act of an interview, more questions could be asked for more information. The interviews went through qualitative analysis whereas the survey's provided a quantitative pool of data.

This research presents its finding from the preliminary research phase, the interviews with experienced software practitioners, and the survey distributed to provide an insight into how the practitioners of the industry uphold their software quality.

Gamification in Educational and Industrial Contexts: Play and Learn (Query Competition)



Edward Langlands University of Canterbury

Supervisor: Dr. Miguel Morales-Trujillo

(Clarendon, 10:00-10:15)

The application of game-based features such as points, leaderboards, and challenges is becoming increasingly popular throughout all areas of education. These features are claimed to positively impact student performance, motivation, and experience. However, the lack of empirical evidence hinders its acceptance as an effective learning and teaching strategy, as the true consequences of its implementation are not widely known.

With the aim of demonstrating the effects of game-based features in a database course, Query Competition (QC) was created. QC is a Structured Query Language (SQL) learning platform that offers students a gamified environment to practice their SQL query skills. The tool allows students to participate in tournaments answering query-based challenges in a competitive setting.

In this project, a systematic literature review was performed, the latest iteration of QC was extended, and an empirical study was conducted. The study was designed to examine the difference in effectiveness each individual game-based element had on students (points and leaderboards). The new version of QC was validated by 131 students enrolled in a database course from the Universidad Nacional Autónoma de México. After analyzing the study results, it was concluded that several aspects of student motivation were improved through the introduction of game elements. No claim could be made about any significant difference the game-based features had on student performance. However, as a whole, participants reported an extremely positive user experience.

Gamified Scrum Retrospectives to Enhance Motivation and Engagement in Learning



Seth Kingsbury University of Canterbury

Supervisor: Dr Miguel Morales-Trujillo

(Clarendon, 10:15-10:30)

Gamification aims to use game-based solutions to facilitate motivation and engagement in non game environments such as education. In software engineering soft skills developed in collaborative environments have become more essential as agile frameworks have been adopted throughout the industry. A key ceremony in many agile development frameworks is the retrospective, an event based around reflecting and improving upon all aspects of collaborative development.

This study aims to facilitate motivation and engagement in the Scrum retrospective process using collaborative gamification techniques. This included researching existing solutions for game-based retrospectives and key game elements that promote collaboration followed by design and implementation of an original retrospective game to be used by students for group software projects.

Implement automated deployment, monitoring and logging to a Docker swarm



Samuel Burtenshaw UC Computer Science Education Research Group

Supervisor: Prof Tim Bell

(Clarendon, 11:00-11:15)

The University of Canterbury Computer Science Education Research (UCCSER) team is responsible for developing tools and software to help further computer science and software engineering studies at schools across the globe. One of the ways UCCSER encourages this is by providing a range of websites that both teachers and their students can use to further their learning in computer science topics both in and out of class. After moving away from google cloud services due to an increased cost and abundance of features that they were not using, the tools were moved to the University of Canterbury Network using containerisation on a docker swarm. This was a much more appropriate funding model and allowed the team to simplify their deployment to only the tools they required.

The aim of this project was to support the development of the University of Canterbury Computer Science Education Research websites by providing an automated deployment pipeline with monitoring and log management. Automation was a key improvement to be made, as previously each update had to be deployed manually. By automating the deployments of tools and software on the swarm the team aimed to improve the efficiency of their deployments and reduce downtime.

This was done by investigating current automation tools for docker and creating an automated tool that could be run periodically on the swarm to search for and identify whether any software was out of date. If any changes had occurred, the tool would proceed to automatically download, update, and redeploy software back onto the swarm. In conjunction with this a custom logging and monitoring stack was deployed to allow the team to view and organise their logs and common metrics as well as be alerted when any errors or issues occur.

Event Registrations for DTHM4Kaiako Web Application



Hayley Krippner UC Computer Science Education Research Group

Supervisor: Prof Tim Bell

(Clarendon, 11:15-11:30)

DTHM4Kaiako is a website designed for Digital Technologies and Hangarau Matihiko (DTHM) curriculum content teachers in New Zealand. The Events section of this Django web application advertises upcoming professional development events for these teachers. The problem space addressed by this project is the need for extending the Events section to allow event organisers to efficiently manage event and event registration data with ease.

A custom solution was desired since a future goal is for registration forms to contain complexly structured custom event session questions. Existing event registration management software does not provide suitable customisation. Key users were interviewed to identify the required core behaviours. A database schema was developed and enhanced to ensure these were captured and modelled correctly. To evaluate the software solution, a think-aloud user study was conducted, where end users participated in walk-throughs to test that the system met the needs of event staff.

Event staff now have an improved data management system for event registrations. Event staff can view their events in the Events Hub. From here, they can view and manage events' data and the data of their associated event registrations. They can track which participants have paid to participate and see their dietary requirements and the counts of each participant type. Event staff can send mass emails to notify event participants of important announcements. Aggregating events and event application data is vital for event staff for a multitude of purposes such as for placing a catering order and processing participants' invoices. Because of this, custom spreadsheets can be generated.

As a teacher, the process for applying for events can be tedious when they must re-enter their details per event registration form. The event registration process is now within the website and participants' data is auto filled in these applications. They can now view a master list of all events they have attended and future events they have applied for. They also can update and withdraw their own event registration forms.

Skill Tracking and Recommendations for CodeWOF



Griffin Baxter UC Computer Science Education Research Group

Supervisor: Prof Tim Bell

(Clarendon, 11:30-11:45)

CodeWOF is a web application developed by the University of Canterbury Computer Science Education Research Group. The app's purpose is to help school teachers maintain and revise their programming knowledge with the use of various coding exercises. Currently, all programming questions on CodeWOF are laid out on a single page, which causes issues when a user chooses a question to attempt. The solution involved modifying CodeWOF to include descriptive tags and filtering for each programming question, tracking the proficiency of users' programming skills, and using this information to suggest exercises for the user to perform. In this modified version of CodeWOF, two question recommendations are provided to users, which are each uniquely generated based on a given user's tracked skills. The first is designed to recommend a question for maintaining programming skills, and the second is designed to be a challenge. As part of this solution, research was conducted involving the teaching, refreshing, tracking, and recommendation of programming skills. After the modifications were completed, the improvements to CodeWOF were evaluated through a survey sent to digital technologies teachers around New Zealand. This evaluation showed that while most users did not find the original questions page to be too bad, almost all survey respondents found the new version of CodeWOF to be significantly better. Notably, users liked the filtering panel, the new questions page with question tags and filtering, and enjoyed the question recommendations.

Recording Design Decisions in Slack



Zachary Kaye University of Canterbury

Supervisor: Dr Fabian Gilson

(Clarendon, 11:45-12:00)

In agile software development, design decisions are often scarcely documented or not documented at all, with working software favoured over comprehensive documentation. However, documented design decisions help share knowledge between developers and help with ongoing maintenance. Research has shown that instant messaging platforms used during software development, such as Slack, contain valuable rationale and design decisions. Therefore in 2019, the idea of recording design decisions in Slack was proposed. Since then, subsequent prototypes of a Slack bot that records design decisions as they are discussed in Slack have been developed. The most recent prototype uses a machine learning model with splitting and classification of complex sentences “roughly” following the WH(Y) template.

In this year’s iteration of the recording design decisions in Slack project, the most recent prototype was further developed before conducting a user study to determine whether the current solution is suitable. The newly implemented features include the storage of recorded design decisions in a database external to Slack, searching and editing stored design decisions, outputting recorded design decisions to a GitLab Wiki page and multi-Slack-workspace support using OAuth 2.0. Although the user study is still ongoing, it has three main objectives. First, to improve the usability of the Slack bot so that the interactions flow naturally between the bot and its users. Second, to improve the natural language processing aspect of the Slack bot that classifies design decisions with users training the bot as they use it. Last, to understand how useful such a tool would be in an ongoing software development project, i.e. would it encourage users to record decisions and what level of quality these decisions would be?

Question tracking system for Moodle



Gongzai Li University of Canterbury

Supervisor: Dr Kourosh Neshatian

(Clarendon, 12:00-12:15)

Moodle is a web-based learning content management system. At the University of Canterbury, Moodle is used as Learn and also as the departmental quiz server. Moodle enables teachers to distribute graded assignments, share documents and lecture notes, and run workshops and quizzes. Moodle has a question bank subsystem using which teachers can create, preview, and edit questions and use them in quizzes or class activities.

The goal of this project was to develop a question tracking system for Moodle's question bank. I have developed a prototype of a revision control system that allows teachers to track all historical versions of a question in the question bank. Using this prototype, teachers can review, restore, or compare past versions of a question. This prototype will be integrated into Moodle in the future.

Measuring bone health features in cortical and trabecular bone with minimal user intervention



Oliver Cranshaw, Michael Woodard MARS Bioimaging

Supervisor: Dr James Atlas

(Clarendon, 13:15-13:45)

MARS Bioimaging (MARS) is a medical imaging company based in Christchurch, New Zealand, and is known for its spectral computed tomography (CT) scanners. The CT scanners, alongside their software, MARS Vision, provide quantitative colour imaging to increase diagnostic information available for clinical use. The scanners output image slices of materials, including a patient's bone, which can then be used to determine a range of bone health features. For these features to be measured, the image slices must be accurately segmented into each bone type.

This project aimed to build a software tooling layer to increase automation and extraction of health features in the cortical and trabecular from medical imaging. The manual selection of these regions of interest (ROI) is very time intensive and prone to human error. Automating this process significantly reduces the time to select the ROI across multiple slices. In addition, the solution generates more accurate data if the neural networks used are trained well. These improvements to the current process will allow clinicians and radiologists to provide more accurate diagnoses of cortical and trabecular bone health.

Another objective of this project was to conduct a thorough analysis of multiple segmentation algorithms. An evaluation was conducted on these segmentation algorithms; this involved a comparison of their measurements of cortical and trabecular health features, along with other metrics such as temporal and spatial complexity. This objective aimed to systematically evaluate each algorithm's advantages and potential shortcomings. Each segmentation method was developed with the intention that once completed, the most effective could be integrated into MARS' bone analysis pipeline.

Retinal image quality assurance using deep learning technique



Nathan Huynh University of Canterbury, Te Whatu Ora Waitaha Canterbury

Supervisor: Dr Andrew Bainbridge-Smith

(Clarendon, 13:45-13:00)

Retinal image quality assurance (QA) is an essential first step for reliably diagnosing macular-related diseases. Deep-learning-assisted techniques have recently gained popularity for classification tasks and are potentially suitable for image QA. We introduced two pre-trained convolutional networks, EfficientnetB4 and VGG16, to classify retinal images into good and bad quality categories. VGG16 was trained on a noisy dataset labelled by a student, and EfficientnetB4 was trained on a small "golden label" dataset labelled by a professional ophthalmologist. VGG16 yielded a higher F1 score of 98%. We only find DRIMDB as the only available retinal dataset labelled with good/bad quality to evaluate the model's performance. However, its image layout is different from our dataset. When evaluating, EfficientnetB4 attained superior performance at 86% F1 score.

We also noticed that there are not many retinal datasets available. A larger dataset with a general retinal image layout and accurately labelled is needed to develop deep-learning-assisted retinal image quality assurance further. We introduced a platform (available on Android and online) for ophthalmologists to contribute their images to the public dataset.

Much like Wikipedia, users can freely upload and label uploaded images. Users can upload their selected images and label or help label other users' images. The platform also hosts our pre-trained model that would predict whether the image is good or bad. This encourages users to contribute to the dataset as an open-source project.

Tumour Region Segmentation in Whole Slide Images



Mitchell Veale University of Canterbury

Supervisor: Prof R. Mukundan

(Clarendon, 14:00-14:15)

In 2020, breast cancer was the most commonly diagnosed type of cancer, making up 11.7% of global cancer diagnoses and 6.9% of global cancer diagnoses resulting in death. An important step in the identification of breast cancer is to extract and analyse a biopsy of the affected tissue. The analysis of a biopsy is usually performed by a pathologist who will view the biopsied tissue under a microscope and flag any regions of the slide that appear to contain tumour related tissue. This manual analysis can be extremely time consuming and has the potential to introduce errors due subjective evaluation that can occur from pathologist to pathologist, and variances between different tumours of the same type, known as tumour heterogeneity.

By using modern high resolution digital microscopes these histology slides can be converted into digital whole slide images (WSIs) which can then be analysed using software, reducing the workload of pathologists, and making the evaluation of WSIs more consistent. This project aims to assist pathologists in analysing whole slide images by offering a machine learning based software algorithm that is capable of automatically identifying tumour tissue region in whole slide images. The segmentation of the tumour regions in WSIs is an important pre-processing step for many machine learning based applications designed for automated classification and grading.

After performing literature review and initial research on tumour segmentation a convolutional neural network approach was developed using the TensorFlow deep learning library. Development of this approach involved designing solutions to efficiently generate large amounts of accurate data that could be used to train the network to differentiate between tumour tissue from other types of tissue. After analysing a WSI the solution is capable of generating an output that can then be used for further analysis or be imported into specialist WSI analysis software to highlight potential tumour regions that were detected.

Preliminary results from the developed network show that it is able to identify tumour regions with $> 85\%$ average accuracy.

3d Reconstructions of Tumour Resections Using Whole Slide Images



Mitchell Freeman University of Canterbury

Supervisor: Prof R. Mukundan

(Clarendon, 14:15-14:30)

Histopathological samples of tissue specimens and surgical resections are traditionally stored in glass slides and examined under a light-field microscope. Recent technological advances in the field of digital pathology have made it possible to obtain high resolution whole slide images (WSIs) of tissue specimens that could be processed using image analysis, machine learning and computer graphics algorithms to provide pathologists and surgeons valuable information useful for diagnosis and treatment.

This project aims to develop and evaluate methods for three-dimensional reconstructions of Ductal Carcinoma In Situ (DCIS) tumour resections from whole slide images of breast conserving surgery specimens. The contours of regions of interested extracted from whole slide images are joined together using splines to fill the gaps between the sparse layers. The reconstructed 3D mesh is then rendered using graphics methods. Different types of splines are evaluated in this project. Kachanek-Bartel splines are found to be particularly useful for 3D shape reconstructions from sparse data. Subjective evaluations of the outputs will be performed by a consultant pathologist.

SESSION: CRYPTOGRAPHY

(time: 15:00-16:00)

Chair: James Atlas

Differential Privacy Tools for Smart Contracts



Kye Oldham University of Canterbury

Supervisor: A/Prof Ben Adams

(Clarendon, 15:00-15:15)

This project aims to implement differential privacy algorithms for smart contracts running on the cryptographically secure Secret Network. Current solutions are created on blockchains with a public ledger, such as Ethereum. User data is exposed on public ledgers, which means there does not exist a differential privacy solution for Smart contracts that is secure. This is why these algorithms will be developed on the Secret Network. To develop a solution, fixed-point arithmetic, pseudo-random number generation and the Laplace algorithm were solved for Secret smart contracts. A library of smart contracts implementing well-known differential privacy algorithms on the Secret Network has been created. They were deployed to the Secret Network's Pulsar test network. Users can view differentially private summary statistics, count and mean, for data stored in the Secret Contracts. This is done with a frontend React application using the SecretJS library. Users may also input data into the contracts using the frontend web application.

Post-Quantum DTLS and MQTT protocols

R

Callum McLoughlin University of Canterbury

Supervisor: Dr Clémentine Gritti

(Clarendon, 15:15-15:30)

Quantum computers are a threat to the current status quo of secure cryptographic communication. It is known that once quantum computers become widely available in the not-so-distant future that the current secure communication protocols will no longer be acceptable. This results in the need to adopt new cryptographic standards, algorithms secure against quantum computing of which there are a few to choose from. One important use case of current cryptography is within Internet of Things (IoT) devices, which generally operate over the Datagram Transport Layer Security (DTLS) protocol. It is crucial that the new standards selected and implemented for a post-quantum version of the DTLS protocol are feasible in terms of the computational, storage, and communication costs involved on these types of devices.

Post-Quantum DTLS and MQTT protocols



Raven Townsend University of Canterbury

Supervisor: Dr Clémentine Gritti

(Clarendon, 15:30-15:45)

MQTT is a common communication protocol used in the Internet of Things. The protocol provides no security, and current proposals to add security do not consider the future risk of post-quantum attacks. This is a severe risk, as MQTT is often used in conjunction with private data and critical systems, such as locks, video cameras and sensors so security breaches could have a significant impact.

The purpose of this project is to investigate potential ways of adding post-quantum-safe authentication security to MQTT.

Ternary Merkle Trees for Use in the Internet of Things



Finn Bright University of Canterbury

Supervisor: Dr Clémentine Gritti

(Clarendon, 15:45-16:00)

The problem addressed by this project is to provide an alternate tree structure for the post-quantum signature schemes known as the Merkle Signature Scheme (MSS) and the Extended Merkle Signature Scheme (XMSS). An Internet of Things (IoT) device would use one of these schemes to create a daily binary tree with which it would sign and send sensor information every hour for 24 messages. Because both schemes use a tree's leaves to sign messages the best accommodating binary tree would need to have 32 leaves. This causes eight leaves to go unused for the process resulting in unnecessary computations for the root node and signature creation.

This project proposes using a ternary tree where all nodes, except for the leaves, would have three children. The best-suited ternary tree would have 27 leaves. Hence, only three leaves are unused and fewer nodes are needed overall for the sensor. Two versions of every scheme were developed where the whole tree was either stored or not. This was done to accelerate signature creation in the former. Results show that both versions of the ternary XMSS algorithms performed better in terms of time and computations than their equivalent binary XMSS algorithms. The MSS algorithms performed faster than the XMSS algorithms but required larger signatures. All of the MSS algorithms performed similarly to each other in time, but the ternary iterations required fewer node hashes. Between the stored and non-stored algorithms, the former needed more memory but performed faster during the signature creation phase.

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.

COPYRIGHT

The cover of this booklet has been designed by Fabian Gilson and is subject to CC BY-NC-ND license, see <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

CREDITS

This booklet has been prepared by Fabian Gilson and has been adapted from the work of Maxime Lucas, Pau Clusella and Thomas Kreuz.

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

Cover image by BiZkettE1 [freepik.com](https://www.freepik.com).



csse@canterbury.ac.nz



@uccsse



uc.csse



uccsse

Te Tari Pūhanga Pūmanawa Rorohiko
Computer Science & Software Engineering

Te Whare Wānanga o Waitaha | University of Canterbury