UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Software Engineering Final Year Projects

Students Posters

2023

Agricultural Land Use Classification

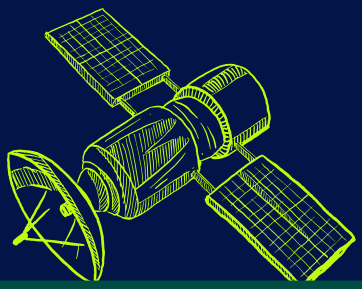

Using Deep Learning

Motivation

Onside Limited is upgrading their biosecurity model, Onside Intelligence (OSI). Onside wants to deviate away from randomly selecting properties and rather create streamlined management method of where biosecurity hazards occur. Currently, the process for adding more resolution to what is on properties is by using "blocks", which are sub-regions within the property which are designated to different crop types.
Onside wants to improve this digital toolkit as it is a manual process, specifically, they want an autonomous method to classify land use and vegetation type.

Solution

- Data provided by Onside in combination with satellite imagery is used to create a training dataset.
- The dataset is used to train a deep learning model that classifies agricultural land use autonomously.

IMPLEMENTATION

A bounding polygon shapefile is used to select the block that wants to be classified

The polygon is divided into areas approximating 1000 m^2

Images are cropped by the individual areas from the divided polygon

Images are shuffled, randomly selected and processed to fit the trained classifier's input requirements

The trained classifier predicts the image and gives it a classification label

If the randomly classified images meet a statistical threshold over the same category then an overall prediction is made

Satellite Imagery

Land Information New Zealand (LINZ) base maps has a web map tile service that provides three spectral bands.
LINZ data was chosen as whilst being open-source and publicly available it also had high resolution imagery ranging from 0.5m - 10m.

"Sourced from the LINZ Data Service and licensed for reuse under the CC BY 4.0 licence."

Results

- Three metrics were used to determine the performance of the model: precision, recall, F1-score.
- The model predicted accurately 90% of the time, and it also has a good balance between recall and F1-score at 0.86 and 0.87, respectively, suggesting it effectively identifies a substantial portion of actual positive cases while maintaining a strong precision.

90% Precision

86% Recall

87% F1-Score

Technology

Python QGIS NumPy

DEVELOPER | ARISH ABALOS

Co-Student | Andrew Cook
Supervisor | Ramakrishnan Mukundan
***Industry Partner* | Onside Limited**

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Computer Science & Software Engineering

State of Practice of DevSecOps in Aotearoa New Zealand

Aotearoa New Zealand, Do you put security in software development and operations?

Only 67% do.

Introduction

Software **security is often neglected** due to its tedious nature. To prove are the recent security attacks across different sectors. **DevSecOps integrates security** into the fast-paced development and operations, so the software industry can meet the market time requirement while producing reliable products.

Objective

Investigate the security practices of the software industry in Aotearoa New Zealand and the challenges they face when implementing them.

Methods

01 Conduct a literature review about DevSecOps and identify the best security practices

"SecDevOps" or "DevSecOps" or "DevOpsSec"

02 Compare the security practices of software professionals in Aotearoa New Zealand to the recommended practices through a survey

03 Conduct a follow-up interview with eager participants for a more in-depth understanding of the respondents' experiences

Future work

Design a security-aware process that software organizations in Aotearoa New Zealand can implement, ensuring that security is considered in every step

Findings

6 out of 15 have **not** gotten any security trainings on their current role

For at least some requirements, security risks are considered when:

73% writing functional requirements

80% designing or writing code

67% implementing secure coding

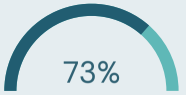

73% monitoring features

Top three most used security tools:

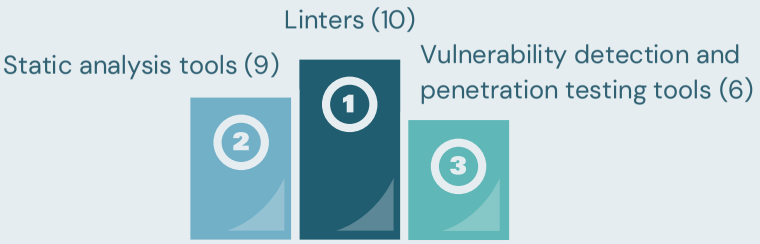

UC University of Canterbury
Computer Science & Software Engineering

Student Darryl Anne Alang

Advisor Fabian Gilson

Machining Chatter Detection with AI

-Lachlan Alsop

UC University of Canterbury Te Whare Wānanga o Waitaha Christchurch New Zealand
Computer Science & Software Engineering

ZEALAND SYSTEMS

Context

Detecting self-excited chatter in CNC milling operations is crucial due to its negative impacts on tools, machines, power consumption, and product quality. Both humans and now Artificial Intelligence (AI) show promise in this task.

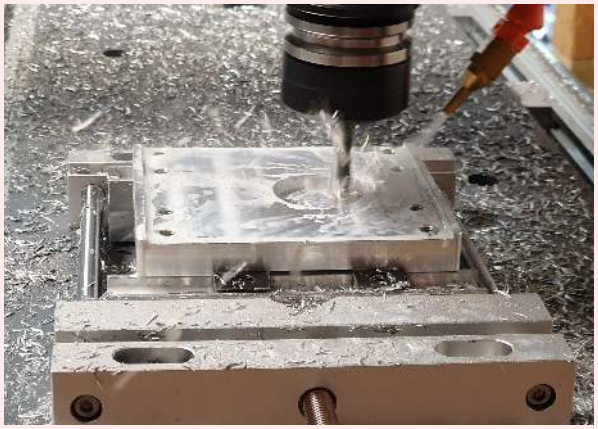

Spectrogram of a milling cut showing the harmonics of milling operation

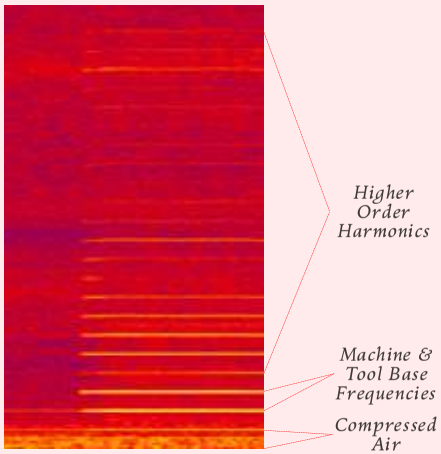

What is Chatter?

Self-excited chatter in CNC milling arises from the high-force interaction between the tool and workpiece. It occurs due to improper setup or aggressive cutting parameters, leading to a natural oscillation induced by the tool's rotation. This oscillation can escalate, causing sustained vibrations, altering the tool's path, and amplifying negative consequences in a positive feedback loop.

Chatter and Sound

Chatter detection's feasibility based on sound relies on two main factors:

- Human Ear as Benchmark: Humans can detect chatter by listening, and since AI excels in signal processing tasks, it should excel at chatter detection.
- Sound Information Availability: Every tool-workpiece interaction produces sound, encoding essential information for detection.

In real workshop settings with background noise like music, compressors, and coolant, AI's resilience is improved by intentionally including noise in the training data.

CNN Processing

To effectively address this challenge on low processing power, like that of an Orange Pi 5, a streamlined approach is needed. Complex AI structures are impractical due to the limited processing capacity. Instead, a basic Convolutional Neural Network (CNN) with minimal layers is employed. This simplicity enhances speed, achieving a processing time of 35ms per second of data, using just 4% of available processing resources. This efficiency allows the Pi to allocate resources to other critical tasks, ensuring a smooth user experience.

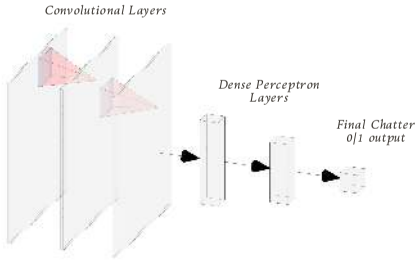

-A simple, but effective CNN structure.

Why is it Useful?

This system not only enables chatter detection and correction but also opens doors to optimizing various aspects of the milling process.

Additionally, it offers the potential for integration with existing research on high-speed milling stability equations. With the developed detection system, these equations can be further refined and optimized.

Moreover, the system can be used to monitor unexpected failures, such as tool breakage, material buildup, or machine crashes.

UC University of Canterbury Te Whare Wānanga o Waitaha Christchurch New Zealand
Computer Science & Software Engineering

EAGLE TECHNOLOGY

esri Official Distributor

City Planning Digital Twin

Purpose

- Utilise the geographic tools made available through the ArcGIS Maps Software Development Kit for JavaScript
- Create a tool to generate digital twins of New Zealand cities using data collected and generated by Eagle Technologies
- Include features that help urban planners visualise changes that they want to make and how it will affect the layout of the city
- Improve city planning by integrating information about important area data. In this case we focused on urban heat islands

Creating Digital Twins

- Create a digital twin of a city as a 3d model using street map data, elevation data, and locations of important features such as trees and buildings
- The data used is supplied by Eagle Technology
- The ArcGIS Maps Software Development Kit for JavaScript is used to collate the data and display it in a web application
- This application allows users to get a detailed view of the city and visualise exactly where things are

City Planning Tools

- The digital twin is expanded with tools to help in planning new developments
- Users can select an area and are able to edit features within that area
- Users can move, update, add and delete features
- Changes are made locally and can be reverted to the original arrangement

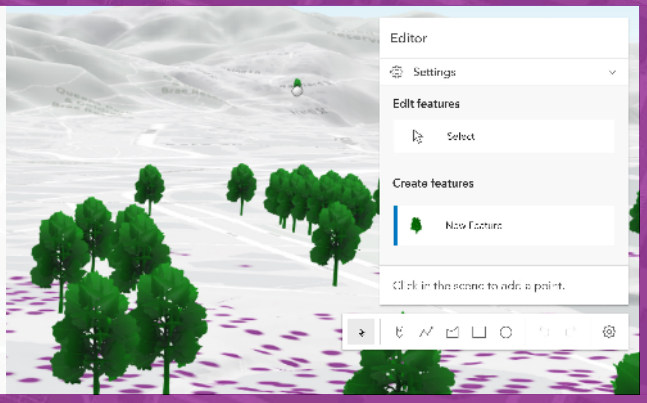

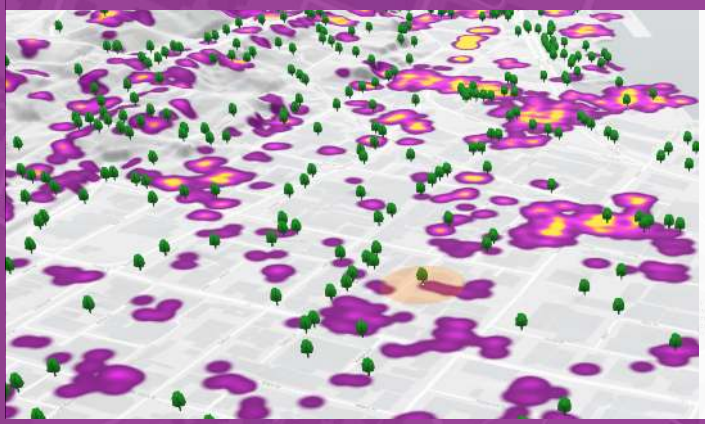

Urban Heat Islands

- Urban structures and materials tend to trap heat within a city, creating urban heat islands
- The urban heat island effect can be reduced and managed with careful placement of buildings and trees
- Modelling heat islands allow us to show urban planners how their planned changes will impact the heat concentration in an area

Student
Max Bastida

Supervisor
Ben Adams

Industry Partner
Edward Wong

Bloodstain Classification Using Deep Learning

By Aidan Campbell, Andrew Bainbridge-Smith, Richard Green, Rosalyn Rough, Oliver Batchelor

Intro

Bloodstain pattern analysis is the practice of investigating bloodstain patterns from crimes and accidents to assist in evidence gathering. This practice is subject to individual bias, experience, and human error and therefore a method is needed to quantitatively analyse a bloodstain pattern. Therefore, this project aims to investigate two objectives:

- How effective are neural networks trained with the image data of individual bloodstains at identifying bloodstain patterns?
- What characteristics of individual bloodstains are critical for neural networks to identify bloodstains?

Methods

- A multilayer perceptron was developed to take the stains' data and output a classification. Information was then removed one at a time to investigate which data was most critical to the network for identification.
- A 3D point cloud network was explored called PointNet++. This network takes the boundary coordinates of all the stains by using their X and Y values and setting their Z values to 1.

Results

Figure 1 displays the result from the multilayer perceptron, and Figure 2 displays the results from Pointnet++.

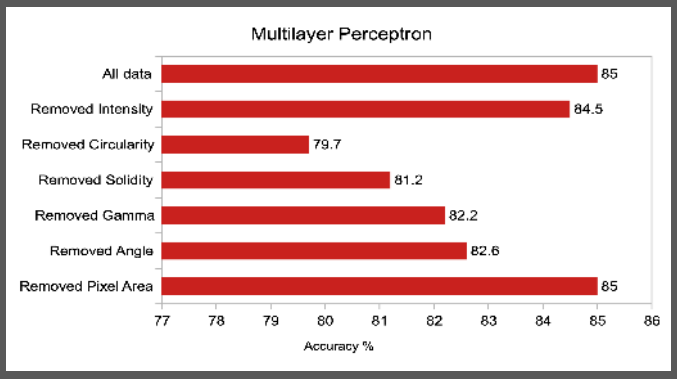

Figure 1: Multilayer perceptron results

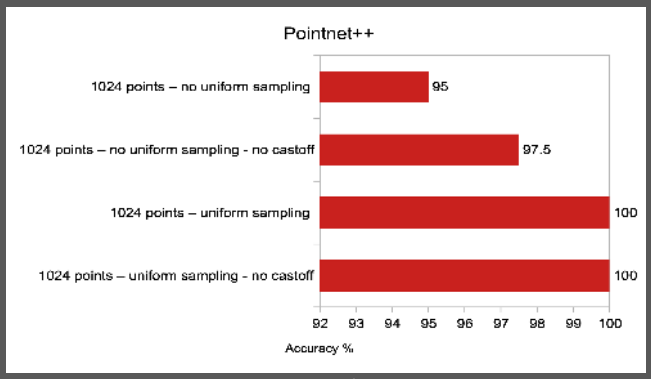

Figure 2: Pointnet++ perceptron results

Figure 3: Pixel coordinates of bloodstain graphed

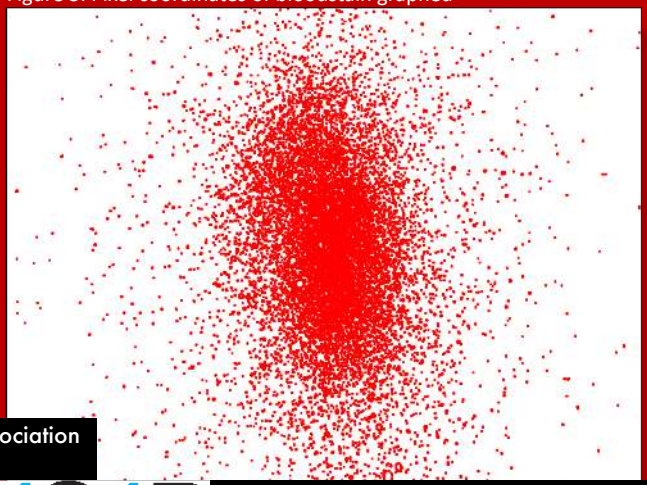

Conclusions

Using a 3D point cloud network such as Pointnet++ proved to be much more accurate than the individual data being passed through a multilayer perceptron. It was also found that circularity of stains was the most important factor in identification.

Pointnet++ achieved 100% accuracy when using uniform sampling, which means it takes the 1024 points that are farthest from each other. This is likely because it captures the spatial information of the entire bloodstain.

However, Pointnet++ is limited in that it is restricted in the number of points it can use. This was set to be 1024 as this was the lowest number of points that was able to capture a reasonable number for all bloodstain types. In future, a network that is able to all the coordinates as input could prove more capable.

In association with:

ESR Science for Communities
He Pūtaiao, He Tāngata

UC University of Canterbury
Computer Science & Software Engineering

Gamification in Educational and Industrial Contexts: Play and Learn

By: George Carr-Smith
Supervisor: Miguel Morales

Introduction

Lectures fall short at keeping students **engaged** and **motivated** while learning. Risk management requires students to identify, analyse, mitigate and respond to risks in a software project, which requires **practice**. We can leverage **gamification** elements such as cooperative tasks, points and joint decision-making to help **motivate** students and effectively teach them risk management.

Challenges

- How do we quantify the effectiveness of gamification in tertiary education?
- How do we keep students motivated in a lecture?
- What is an effective way to teach students risk management?

Objectives

"How do collaborative game elements in a game-based solution increase levels of student motivation, engagement and satisfaction in a course at a tertiary education institution?"

We answered the research question by creating a game based on Pictionary called **Risktionary**. We used the **action feedback loop** to complete the following objectives:

- Develop the Risktionary game
- Evaluate survey results to determine student motivation, engagement, and satisfaction
- Evaluate survey **feedback** to plan improvements
- Iteratively **improve** the game based on feedback

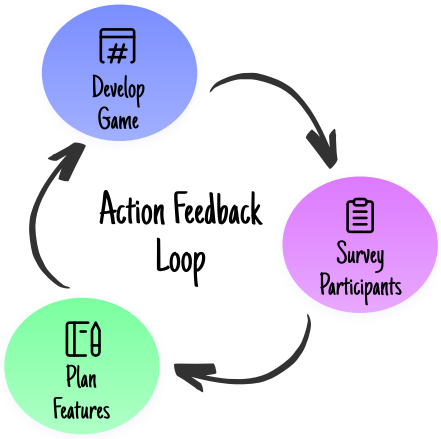

Play A Round of Risktionary

1. Draw or Guess

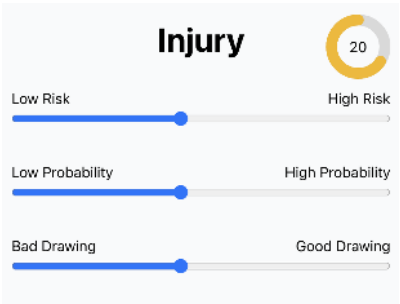

2. Vote on Probability and Severity

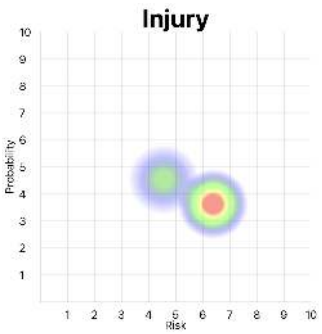

3. Discuss Votes

4. Learn more about the risk

Results

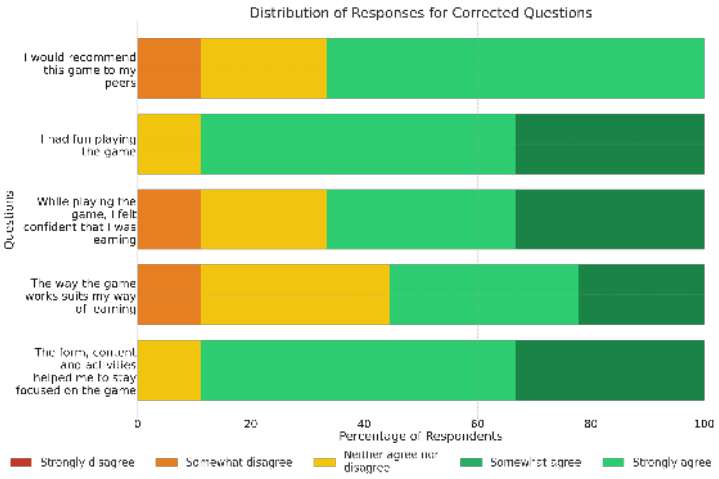

Conclusions

- Most participants agreed that the game was fun, helped them feel confident that they were learning and kept them focused.
- Students voiced their appreciation for the game, with one remarking, "I had fun guessing the words! It's like skribbl.io but for risks, definitely an engaging way to learn."
- The game was effective at increasing student motivation, engagement and satisfaction.
- Gamification stands out as an effective method for teaching students risk management.

UC University of Canterbury
Computer Science & Software Engineering

React Node.js Socket.io

Branch Cut-Point Detection

Manual Pruning

Time-consuming

Hazardous

Objectives

Automatic pruning by a drone requires accurate **detection** of the **cut-point**.

- **Train** a neural network to detect the cut point.
- **Track** the cut point over a sequence of frames.

Solution

R-CNN Deep Neural Network trained on 123 annotated images to detect cut-point region.
An Unscented Kalman Filter tracks the centre.

OpenCV Detectron2

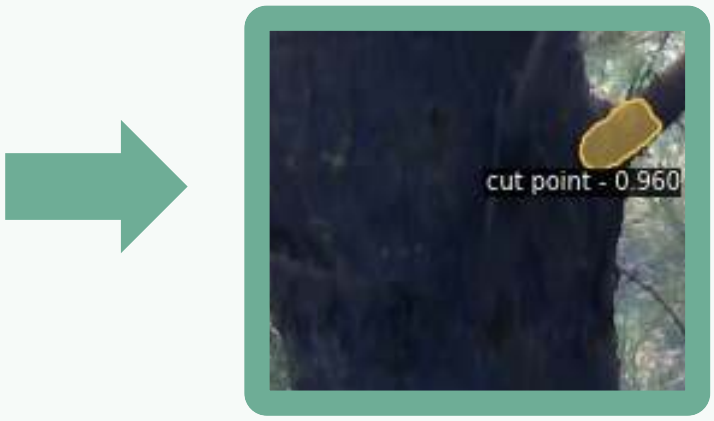

Results

- Average Error: **88.63** pixels.
- Detection Rate: **89.32 %**.

The detector is most effective when there is **minimal movement**, and at a distance of approximately **1m** from the branch.

Future Research

- Account for excessive motion.
- Account for frames with no detection from the neural network model.
- Obtain 3D co-ordinates to cut-point.

UC University of Canterbury | Computer Science & Software Engineering

MINISTRY OF BUSINESS, INNOVATION & EMPLOYMENT
HIKINA WHAKATUTUKI

Fa Wren Chong

Supervisors
Sam Schofield
Richard Green

Special Thanks
Bradley Scott

UC UNIVERSITY OF CANTERBURY
Computer Science & Software Engineering

Student: James Claridge
Supervisors: Professor Richard Green, Dr Sam Schofield

Branch Detection For Drones

BACKGROUND

- Forestry industry looking to use drones to automate hazardous jobs.
- Branches near powerlines need to be pruned to prevent power outages and fires.
- Pruning branches near powerlines is a dangerous task that risks injury and death.
- Work underway to develop drones to prune branches, but **the drones need to able to identify branches.**

OBJECTIVES

- Develop a system to **identify branches.**
- The system should **label all pixels** belonging to an identified branch.
- The labelling of branches should be **accurate and precise.**
- Branches from different **species of trees** should be recognized by the system.

FUNDING

This project is funded by MBIE as part of the larger project UOCX2104 –" Enabling unmanned aerial vehicles (drones) to use tools in complex dynamic environments"

MINISTRY OF BUSINESS, INNOVATION & EMPLOYMENT
HĪKINA WHAKATUTUKI

SOLUTION

- **Deep Learning** allows a system to recognise an object in varying environments.
- **Pre-Existing** networks exist online that have been trained on vast image databases.
- **Fine-Tuning** enables utilising the knowledge of a pre-existing network while training it to detect a new class of object.
- By using fine-tuning, a pre-existing network can be **re-trained to detect branches.**.

METHOD

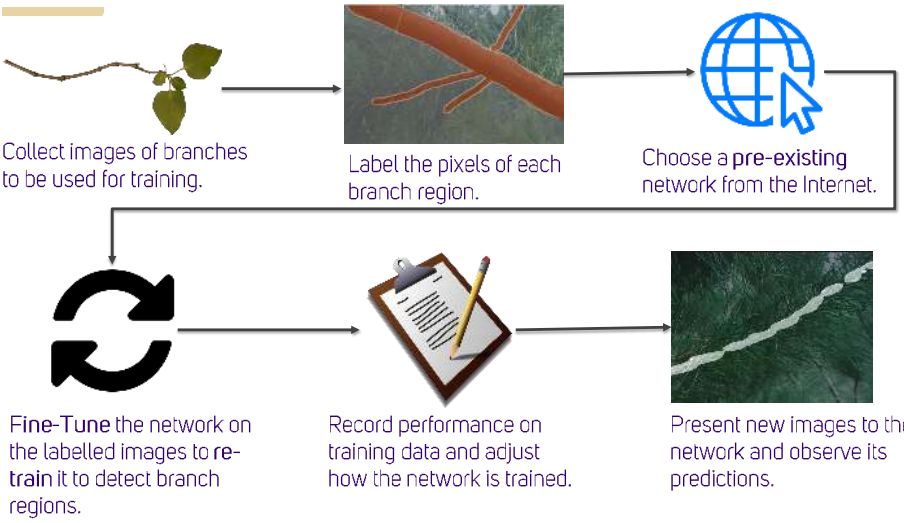

IMPLEMENTATION

- **Semantic Segmentation** networks label regions in images as belonging to a class.
- Each image in the training data was portioned into regions labelled either "Branch" or "Not Branch".
- A pretrained version of the **Segformer** network was fine-tuned on the training data.
- The network was first trained on small images containing only one branch, then on larger images with multiple branches.

RESULTS

Image presented to network.

The networks prediction in green.

EVALUATION

- The final network had an **accuracy of 72.69%** on its training data.
- The network was able to **accurately recognise and label branches in new images** it had not seen before.
- The accuracy could be improved by **increasing the size of the training dataset.**
- While the network performs well on a desktop computer it is likely **too slow to work well on a drone.**
- **Future work** would involve testing the system's performance in the field.

Explainable AI in an Intelligent Tutoring System

PROBLEM

Intelligent Tutoring Systems are learning environments that adapt based on students' wants, needs, and preferences

Students **learn more when they accept recommendations** from the system.

However, students **often do not accept recommendations**, possibly **due to a lack of trust**.

SOLUTION

Adding **XAI may increase trust** and therefore students' learning.

We added explanations to answer the questions:

- **Why was this problem recommended?**
- **Why does problem selection matter?**
- **How is my student level calculated?**

EVALUATION

We studied the effects on 15 students from COSC265, a relational databases course for undergraduate students.

Participants:

- took a personality test
- used the system with an eye tracker
- answered a questionnaire about the explanations

RESULTS

Summary:

Participants	**15**
Time (mins)	**33.7 (3.5)**
Attempts	**9.5 (3.4)**
Solves	**8.7 (4.3)**
Feedback	**72.3 (41.9)**
Used constraints	**99.9 (24.6)**

Students commented...

- they **trusted the system more** once they knew why it was making decisions
- they wanted **more concrete steps** for applying their knowledge
- they would sometimes prefer **explanations to be combined** when they are related concepts

Key takeaways:

- More conscientious and agreeable students view explanations less
- More conscientious students solve more problems
- More open students accept recommendations from the system more, regardless of whether they view explanations

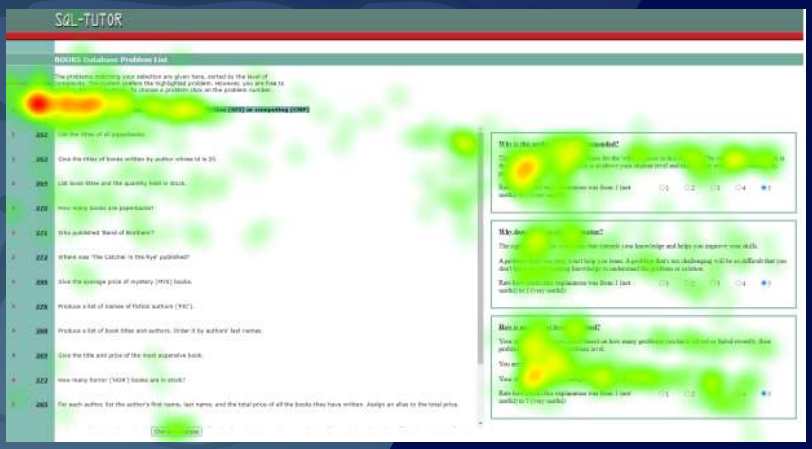

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Tanja Mitrovic and April Clarke

Tumour Segmentation

Finding cancer regions in breast cancer whole slide images using Machine Learning

Sam Clark & Ramakrishnan Mukundan

Breast cancer is a widespread and serious disease that affects millions of people worldwide. Detecting it early and diagnosing it accurately are crucial for effective treatment and improving survival rates. Currently, pathologists examine images of microscopic breast tissue samples to classify the cancer and plan treatment. However, this manual analysis is challenging and can yield varying results. **Computational pathology** offers a solution by using **computer vision** to analyze large images of tissue samples, with more consistent results.

We create a **machine learning model** to **segment tumour regions** in breast cancer **whole slide images**. Most existing Research in this area uses complex deep learning models to segment tumour regions. This can be done highly accurately, however the justification as to 'why' the model makes it's decisions is much less understandable to humans. Thus, our research seeks to create a segmentation model that is both **accurate and interpretable**.

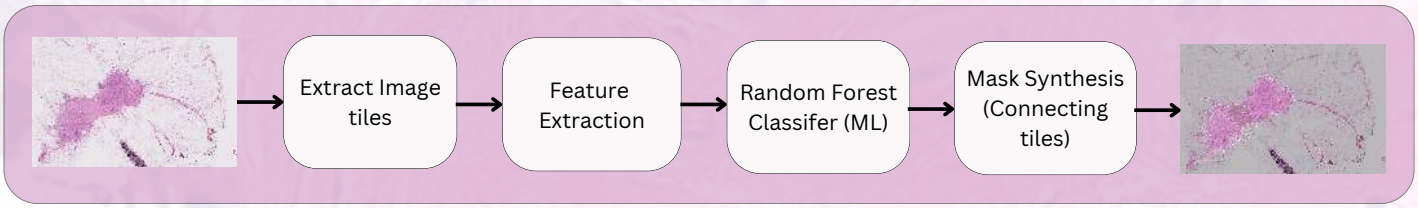

Method

We use a **divide and conquer** approach to segment tumour regions. Firstly, we **break** the **whole slide image** into **tiles**. Then we extract selectively chosen **texture features** using algorithms such as Grey level coocurence matrices and local binary pattern. With these features we create a feature vector for each tile which is passed to a random forest **machine learning** model to **classify each tile** as either cancerous or not. Then we **synthesise** the tiles back together to form a binary cancer **segmentation mask**.

Results

We created a machine learning model that segments tumour regions with a median accuracy of 93.5% and a median F1-score of 78.7%.

Median Accuracy: 93.5%
Median F1-score: 78.7%

It was determined that mean **hue** and **saturation** components of tiles, and the correlation or **Grey-level coocurence matrices** are the most **prevalent fetures** for classifying tumour tiles.

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Agricultural Land Delineation using Deep Learning

Andrew Cook

Motivation

- Our industry partner has a biosecurity model allowing for biosecurity threats to be traced throughout the country
- For this model to work, a significant amount of data from users is required. Specifically:
  - Who is visiting which properties, and when
  - What is being grown on a property
- Automating the process for adding this data allows for easier customer use, and more data for the biosecurity model
- Existing properties can be backfilled with this new generated data

Process

Data

Using QGIS on Toitū Te Whenua 's* Basemap data, 2.64GB of training images were created using data from our industry partner, alongside their true boundaries as a mask

*www.linz.govt.nz/linz-copyright

Training

- Using PyTorch and Python, utilising DeepLabv3 as the architecture for Deep Learning
- A dataset consisting of images and respective masks:
  - 70% Training data
  - 15% Validation Data
  - 15% Testing Data
- The model takes an image, uses binary segmentation to predict what the boundary should be, and compares it to the true boundary to learn how to do the next image more accurately

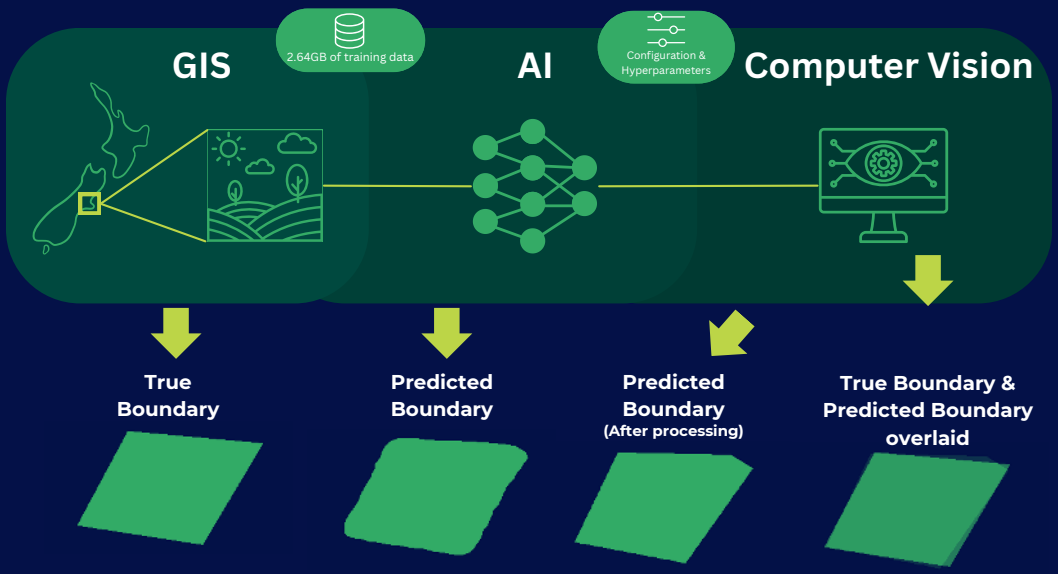

Accuracy

- Used to determine how 'good' the model trained is
- Calculates a score by taking all of the pixels that are shared between two masks, and divides it by all of the pixels if you put the prediction and truth on top of each other
- A loss in accuracy can be taken to simplify the boundary significantly whilst semantically remaining the same

$$\frac{I}{U} = IoU$$

Accuracy of result:

>85%

Co-Student | Arish Abalos

Supervisor | Ramakrishnan Mukundan

UC University of Canterbury
Te Whare Wānanga o Waitaha
Christchurch New Zealand

Computer Science & Software Engineering

BIOSECURITY APP FOR EXOTIC SPECIES REPORTING:

P.I.P

PORT IDENTIFICATION FOR PESTS

Protecting the natural environment and primary industries of Aotearoa New Zealand is important for our unique biodiversity and economy. Exotic species spread via trade routes, exacerbated by climate change that assist spread and intensify impacts. Stopping invasive pests and diseases at the border such as seaports, removes threats and protects the country

OUR MISSION

PIP (Port Identification for Pests) aims to provide port operators with a tool to assist in preventing exotic pests from breaching the border. We will do this by providing a smart app to allow reporting and identification of suspicious biological organisms on sea freight, along with ability to alert port staff to new threats. The outcome will be that the biosecurity risk posed by the sea freight pathway will be significantly reduced

1. Information about a variety of species are shown and sorted for users to help identification for people without expert knowledge.
2. Automated collection of time, location, and shipment numbers speeds up the process, encourging use in a fast-paced port environment.
3. The database of pests is hosted by a host server, which allows them to continually update the information.
4. Being developed in Flutter means the application can be installed and used by Android and IOS (Apple) phones.
5. Integrating Google services allows the use of location and map tools, which streamlines the report and ID process.
6. Surveying port operators for prototype feedback allows the app to be refined specifcally for port-operator users.

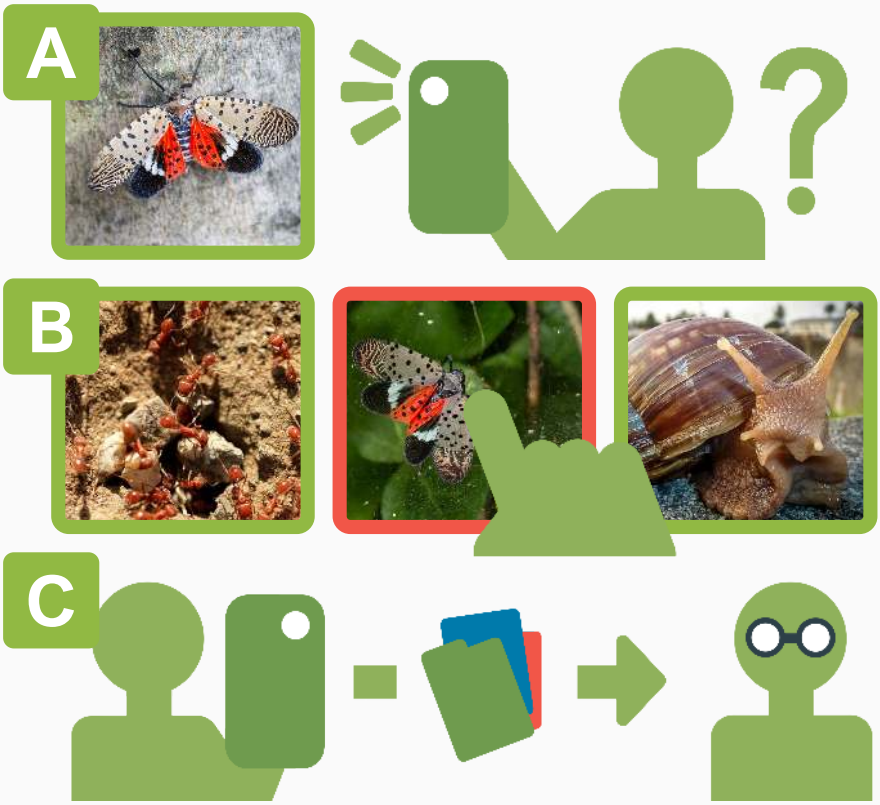

LIVE APPLICATION EXAMPLES

HOW IT WORKS

A. When a port user encounters an organism on sea freight, they take a picture via the app along with information on type of freight and/or sea container ID.
B. The internal database will sort through known species and provide a list of likely pests which allows the user to identify the type of organism (e.g., hitchhiker ants) or the species (e.g., lanternfly).
C. Once selected, this information as well as encounter time, location, freight type and history, is sent as a report to a hosting site for collation onto a database.

Push messages can be sent to app users alerting them of potential risk organisms arriving on sea freight, based on port of origin or seasonality of key pests.

MILESTONES

A. Interactive pest information database.
B. Easy to analyze report database, hosted externally.
C. Image attatchment to improve reports.
D. Prototype database featuring real data on pests.

A Collaborative Project Between

University of Canterbury Computer Science & Software Engineering

agresearch — āta mātai, mātai whetū

Developed By: Lane Edwards-Brown
Supervised By: James Atlas

Self-Regulated Learning in EER-Tutor

Student Sami Elmadani

Supervisor Tanja Mitrovic

Motivation

The motivation behind this study is to improve students' learning experiences by enhancing their self-regulated learning (SRL) skills. By integrating SRL strategies, we aim to optimize student performance and problem-solving abilities when using EER-Tutor.

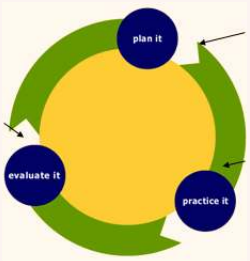

Zimmerman's Model of SRL

Objectives

This study builds on previous work with SQL-Tutor and aims to integrate **Zimmerman's model** into EER-Tutor. Answering the following questions:

RQ1: How would adding self-regulated support using the Zimmerman model affect a student's learning using an Intelligent Tutoring System like EER-Tutor?

RQ2: How would a student's self-regulated learning skills be improved by using EER-Tutor with self-regulated support?

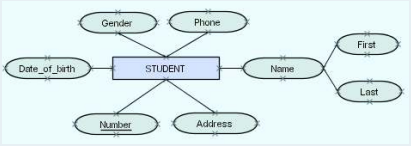

Problem Selection

- In the **self-reflection** and next **forethought phase**, students select the next problem to solve.
- Options include
  - Self-selection
  - Sequential
  - Similar in size
  - Goal concept-focused
  - Related to a new concept
- Encourages continuous reflection and planning. Especially using the **Student Model**.

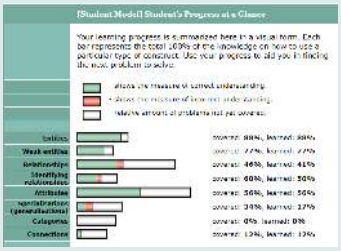

Self Assessment

- Before each problem, students perform a pre-problem self-assessment.
- This assesses their readiness and strategies for solving the problem.
- After each problem, a post-problem self-assessment reflects on performance and guides future adaptations.

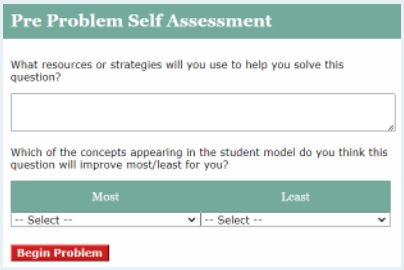

Goal Setting

- In the **forethought phase**, students set specific, challenging, and attainable goals for their learning
- Goals include session duration, problem-solving targets, and database design focus.
- The purpose is to support self-regulation by facilitating goal-oriented learning.

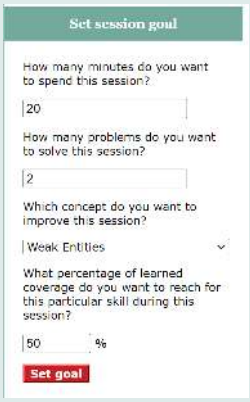

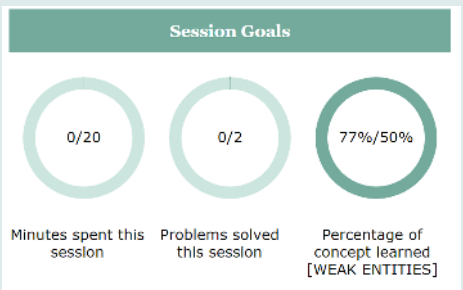

Results

General System Data Summary: These metrics offer valuable insights into how students interacted with the EER-Tutor system and how they performed throughout the self-regulated support intervention.

	Min	*Max*	*Median*	*Mean (SD)*
Number Constraints Used	18	195	139	142.16 (22.9)
Problems Solved	0	35	16	15.43 (9.85)
Feedback Messages Seen	10	656	184	220.24 (158.11)
Logins	3	36	15	15.85 (7.49)
Total Time	12	1624	477	509.88 (271.66)
Attempted Problems	4	47	20	21.31 (10.40)
Number of Attempts	10	323	96	110.81 (74.41)
Pre-Test	2	7	5	4.46 (1.26)
Post Test	0	7	5	5.12 (1.76)
SRL Pre Instrument	2.08	4.83	3.15	3.26 (0.55)
SRL Post Instrument	1.17	4.83	3	3.23 (0.88)

Pre-Assessment: What resources or strategies will you use to help you solve this question?	Use of Lecture Notes and Study Materials
	Utilizing personal knowledge
	Problem-solving approaches
	Seeking Help or Clarification:
	Breaking down the problem into components
	Miscellaneous Responses (uncertainty)
Post-Assessment: What did you learn from solving this problem that you can apply to future problems? Consider your session goals	Learning and Knowledge Gain: Already knew the solution
	User Interface and Software Usage
	Read question carefully.
	Confusion and Ambiguity

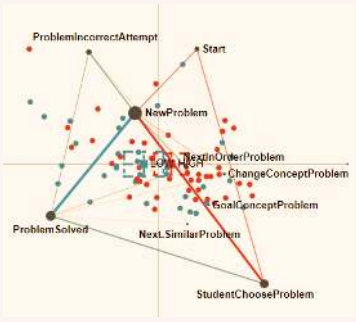

Self-Assessment Responses: Common student responses from both the pre and post-assessment phases of our study. It provides insights into students' problem-solving strategies, resources, and reflections on their learning experiences.

ENA Results: This ENA diagram shows how different students approach learning based on their performance levels. It reveals that low-group students (who scored below 3.26 on the SRL instrument) who choose problems themselves tend to do better, while high-scoring (above 3.26) students rely more on the system's choices.

UC University of Canterbury Te Whare Wānanga o Waitaha Computer Science & Software Engineering

In Collaboration With:

Visit The Website Yourself!

BARRACUDA: A NEW CODING LANGUAGE

FOR ARBITRARY CODE EXECUTION ON GPUS

MOTIVATION

- Many programming languages have features to **execute arbitrary code at runtime**.
- However, **no such feature exists for GPU languages**.
- Barracuda is a new programming language that **fills this gap**.

OBJECTIVES

- **Fix** the existing compiler that doesn't work.
- **Add three new features** to the language: **pointers, arrays, and a type system**.
- **Optimise** the compiler

HOW IT WORKS

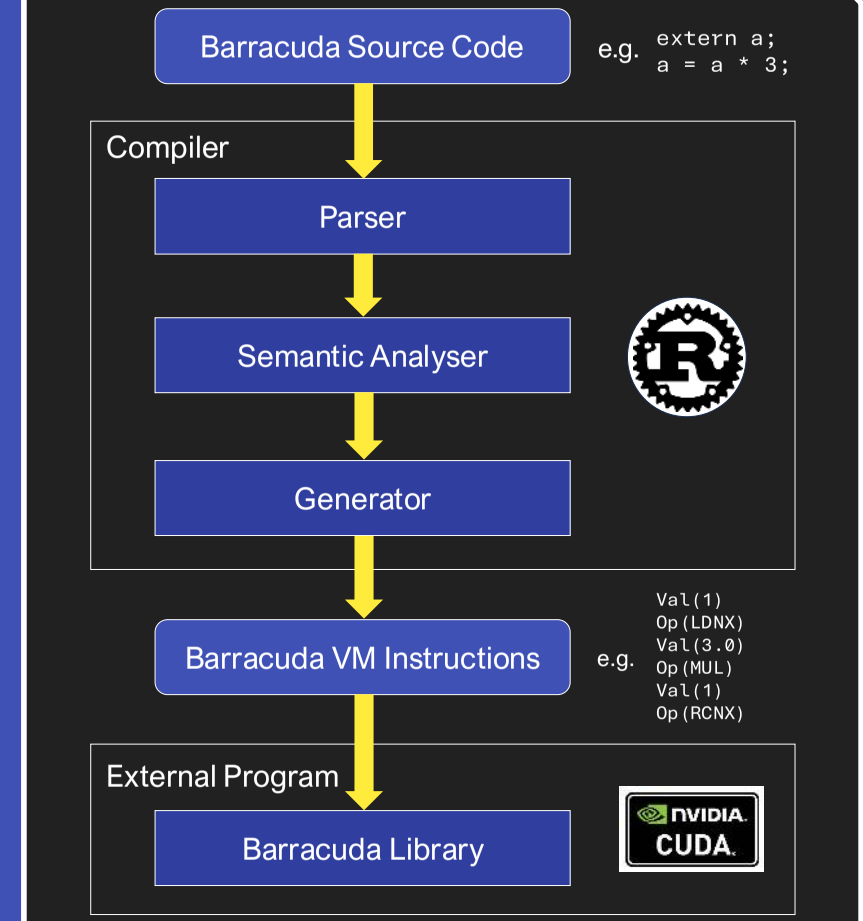

PROGRESS

- Added the **new features** to Barracuda.
- Used the type system to allow for multiple dispatch in functions.
- Wrote two **test suites** to increase confidence and find bugs.
- **Fixed bugs** found during testing.
- **Investigated potential optimizations** for the compiler.

EXAMPLE

```
fn populate_array(array: *[i32; 10],
i, multiplier) {
   *array[i] = *array[i]*multiplier;
}
let array: [i32; 10];
for (let i = 0; i < 10; i = i + 1) {
   populate_array(&array, i, 3);
}
let total: i32 = 0;
for (let j = 0; j < 10; j = j + 1) {
   total = total + array[j];
}
array = [0,0,0,0,0,0,0,0,0,0];
```


Student **Luke Garside**

Supervisor **Fabian Gilson**

In collaboration with **Phillip Duncan** and the **UC Physics Department**

UC UNIVERSITY OF CANTERBURY

Computer Science & Software Engineering

{code:WOF} An Exercise Creation Interface

What Teachers Gain by Asking Questions

George Hampton
Supervised by Tim Bell

Context

In schools, teachers may only engage with programming once a year, and thus forget details in the months between.

CodeWOF is a website enabling teachers to maintain their programming skills by completing simple Python exercises.

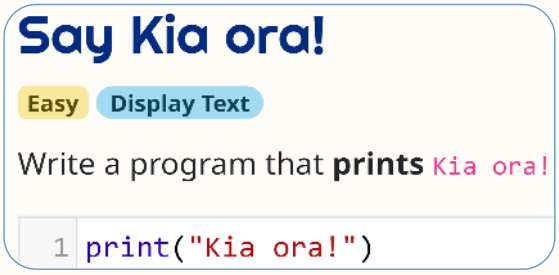

An example CodeWOF exercise

Objectives

The existing method of adding new exercises involved manually creating a technical file describing it, which is error-prone and inconvenient. Thus, the following steps were taken:

- Design and implement an interface for adding CodeWOF exercises.
- Investigate possible benefits of allowing teachers to create exercises using this interface.
- Evaluate how CodeWOF might be adapted to consider AI tools.

The question creation interface, with the form for defining it (left) and the question preview (right). The callouts on the right show substitution of different values for the macros defined in the form.

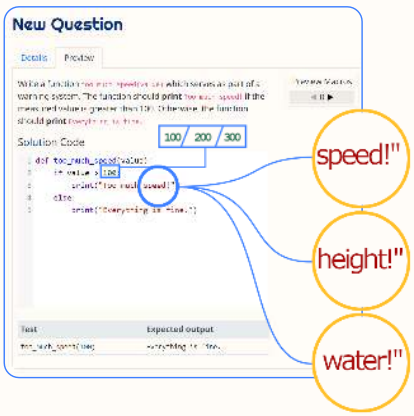

Solution

The interface shown above was designed and implemented with the existing Django technology.

The single-page design of the form presents users with a condensed way to create new questions and their variants, as well as their test cases. Users are then able to submit questions for moderation from a second page.

A qualitative study was conducted to identify teachers' response to the new interface and how it may benefit them, in the form of a questionnaire.

Results

Teachers' responses to the survey detailing potential benefits or drawbacks of the new interface are compiled in the word cloud to the right.

Based on this information, the interface was made available to teachers.

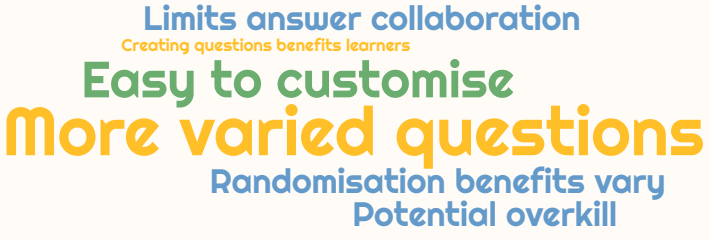

A compilation of comments on benefits from surveyed teachers. The size of the comment is proportional to the number of times it was made.

codewof.co.nz

UC UNIVERSITY OF CANTERBURY
Te Whare Wananga o Waitaha
CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Patient Deterioration Detection

Background

In 2022, Christchurch and Middlemore Hospital were overwhelmed with patients leading to **2 deaths** in busy waiting rooms. The objective of this project is to reduce the number of patients suffering undetected health deteriorations and improve patient health outcomes.

Solution

A **non-contact, camera-based** system to monitor patient **movement** and **body-temperature**. The system prints an alert when a fever is detected in a patient and when a patient is not moving (including not breathing). The system uses a colour camera for motion detection and a thermal camera for fever detection.



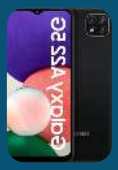
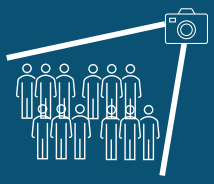

Algorithms

Motion Detection

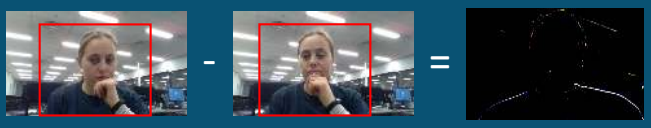

Subtract the current bounding box from the previous and evaluate the remaining image to identify motion

Fever Detection

Identify the temperature of objects in the image and alert if the temperature > 38.3C

Results & Conclusions

The results of testing this life saving camera system determined that motion detection had 90% accuracy and the fever detection had 92% accuracy.

	Precision	Accuracy	Recall
Motion Detection	100%	90%	80%
Fever Detection	83%	92%	100%

This system is a unique idea that would allow patient deterioration to be detected remotely. The system alerts medical staff if a patient has a fever or if a patient is no longer moving, allowing crucial immediate intervention.

This system could save lives.

UC University of Canterbury
Computer Science & Software Engineering

Student
Rachel Hodgson

Supervisor
James Atlas

CODING WITH CONFIDENCE: THE ROLE OF TOOLS IN CODE REVIEW EXCENLLENCE

STUDENT – DANISH KHURSHEED JAHANGIR

SUPERVISOR – DR. FABIAN GILSON

INTRODUCTION

Code reviews are an essential part of software development where developers can examine and evaluate each other's code changes. Our research project is dedicated to identifying and implementing best practices and tools that can enhance both the effectiveness and efficiency of code reviews.

METHODOLOGY

For the purpose of our research, we first compiled information from various research papers on the topics of code review tools and code review best practices. We gathered these papers from websites such as Google Scholar, IEEEXplore, ACM Library, UC library and more.

After conducting the literature review and identifying common tools and best practices, we conducted a focus group discussion with 3rd Year Software engineering students who were encouraged to use certain tools to gain insight into their experiences, methods and impacts of those tools in their code review processes.

RESEARCH QUESTIONS

The research questions that our study aims to answer are:

- What are the common tools and platforms used to perform code reviews?
- What are some of the best practices used in performing code reviews?

Based on the results from the above two questions and the focus group discussion, 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.

COMMON TOOLS

Some of the common tools and platforms found that were used in the code review process were:

- FindBugs
- PMD
- CheckStyle
- CppCheck
- Gerrit
- Github
- Review Bots
- SonaQube

BEST PRACTICES

Some of the best practices used in code review processes are:

- Automating Code Review Tasks
- Establishing a Review Policy
- Providing Clear Feedback
- Understanding the Context and Code
- Preparing the Code Change for Review

CONCLUSION

By leveraging the common tools and platforms found, along with the best practices, we can increase the efficiency and effectiveness of the code review process in our projects. We can use the tools to automate tedious tasks and adhere to the best practices found to deliver high-quality code and reduce the likelihood of defects.

Related literature

- L. MacLeod, M. Greiler, M.-A. Storey, C. Bird, and J. Czerwonka, "Code reviewing in the trenches: Challenges and best practices," IEEE Software, vol. 35, no. 4, pp. 34–42, 2017.
- R. Tufano, L. Pascarella, M. Tufano, D. Poshyvanyk, and G. Bavota, "Towards automating code review activities," in 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE), IEEE, May 2021.
- A. Bacchelli and C. Bird, "Expectations, outcomes, and challenges of modern code review," in 2013 35th International Conference on Software Engineering (ICSE), IEEE, May 2013.
- V. Balachandran, "Reducing human effort and improving quality in peer code reviews using automatic static analysis and reviewer recommendation," in 2013 35th International Conference on Software Engineering (ICSE), IEEE, May 2013.
- D. Singh, V. R. Sekar, K. T. Stolee, and B. Johnson, "Evaluating how static analysis tools can reduce code review effort," in 2017 IEEE symposium on visual languages and human-centric computing (VL/HCC), pp. 101–105, IEEE, 2017.

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha

Computer Science & Software Engineering

CONSTANT CONNECT

Memory Assistive Technology for Dementia Patients

Cody Larsen, in partnership with Hannes de Bruin, supervised by Fabian Gilson

Background

Dementia is a broad term for a set of different disorders that affect the brain. Patients with dementia often experience changes in their mood and behaviour. It is also a progressive disease, meaning that it gets progressively worse as more cells in the brain die. This is a very taxing health issue for patients, and it also is very burdensome for their family members and caregivers. Some of the struggles faced by caregivers and family members when caring for a patient with dementia can be; agitation or aggression, repetitive speech or actions, wandering or restlessness, and much more.

Fig 1 – Prototype for Constant Connect Product

Features

Constant Connect makes use of phrase recognition technology to automatically detect phrases taken from a predefined set of commonly spoken phrases by patients with dementia. The detected phrase will then be used to show a relevant pre-recorded response from a family member. This use of pre-recorded responses and phrase recognition will help to reduce the load placed on family members and caregivers and help to assist the patients in an effective and personal manner.

The solution also aids dementia patients by modelling 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.

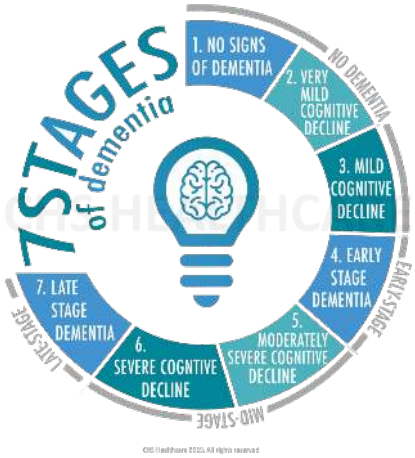

Solution

"Constant Connect" is an innovative assistive technology that aids dementia patients by simulating real conversations with family members. It employs advanced phrase 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.

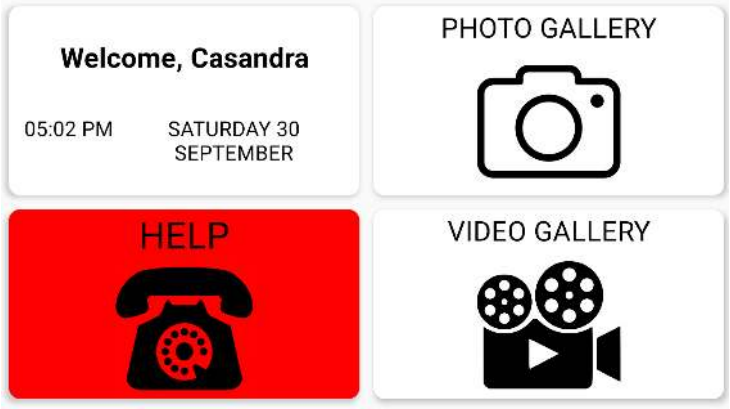

Fig 2 – Final UI for home page of Constant Connect

Evaluation

Constant Connect allows for significance ease of use that will improve the overall quality of life for dementia patients and their family. It successfully bridges the gap between low burden on family and meaningful connections, offering patients the gift of simulated real conversations with their loved ones.

Phrases spoken by the patients are successfully detected by the phrase recognition technology, and automatically supply a pre-recorded response from the family member to seamlessly simulate real conversation for the dementia patient.

The implementation of the photo and video galleries allow for immersive patient care through reminiscence, also allowing loved ones to send their best wishes and memories for the patient to view at any time. Additionally, intermittent and scheduled reminders further enhance the effectiveness of the product and effectively supply essential patient care ensuring the patient is cared for.

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Smooth Drone Formation Flight: Fast Formation Configuration

Frederik Markwell

Supervisor:
Andreas Willig

Team Members:
Euan Morgan, Ethan Ng, John-Paul Lay, Joshua Ellingham

Motivation:

Drones can be used in formation to solve a number of real world issues. Possible applications of this technology include search and rescue, tracking of flying insects and firefighting.

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.

A solution was developed to simplify the process of transferring software to the drones.

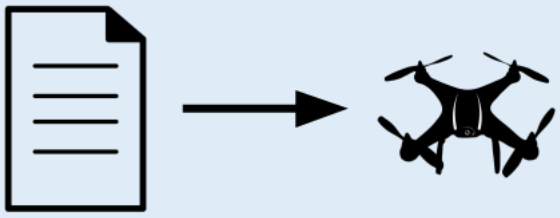

Transfer of Software to Drones in the Field

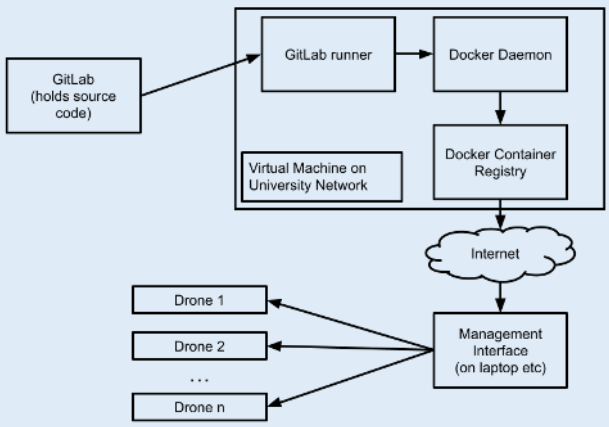

Solution:

The developed solution uses Docker along with conventional Linux tools such as scp to transfer software and configuration files.

Most software deployment systems assume that there is a constant internet connection. When working in the field, we may not have that luxury. We use a two step process, where you can transfer build code to a laptop when you have internet connectivity, and then in the field deploy that on the drones.

The backend was developed in Python, and the frontend in JavaScript using ReactJS.

Transfer Process (from laptop to drone):

1) Detect drones using ARP broadcast and HTTP request to each drone
2) Transfer TLS certificate
3) Synchronise drone clock to laptop's
4) Host Docker registry on laptop
5) SSH into drone to run `docker pull`

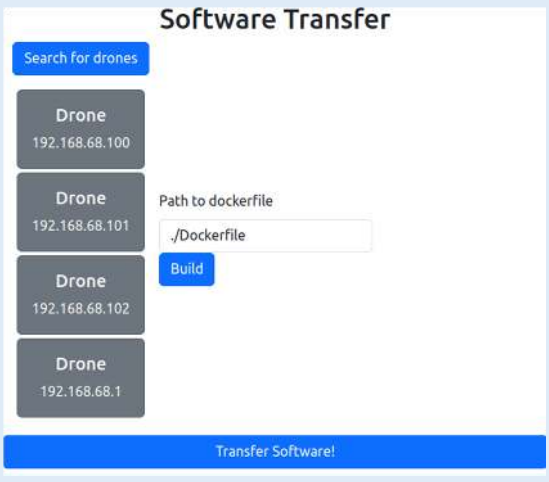

UC UNIVERSITY OF CANTERBURY
Computer Science & Software Engineering

Mixed Reality Collaboration in Geoscience

Using HoloLens Devices to Share Geology

Jack McCorkindale

Academic Supervisor: Matthias Galster
Industry Partners: Tim Schurr, Adam Hunt

MADE WITH Unity · Microsoft HoloLens · Visual Studio

Background and Motivation

- Understanding the geology of a site is indispensable across industries
  - Construction
  - Mining
  - Environmental management.
- Seequent's existing tool, Seequent Central, allows users to create 3D models using geological data
- Consulting geoscientists require user-friendly tools that allow their clients to navigate scene effortlessly
- Mixed Reality is becoming more valuable as the technology develops
- Previous prototypes demonstrated weaknesses that impacted their use

Opuha Dam

Objectives

- Enable the projection of 3D models of geological scenes extracted from Seequent Central
- Facilitate simultaneous viewing of these models by multiple users to improve
  - Collaboration
  - Shared understanding
- Design an intuitive user interface and control system within the solution
  - Minimize the learning curve
  - Ensure users can easily navigate and interact with models

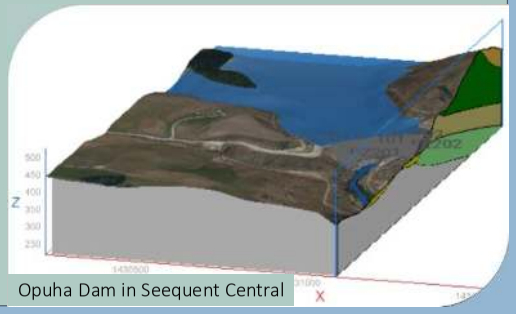

Opuha Dam in Seequent Central

Solution

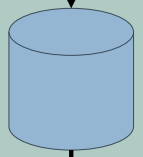

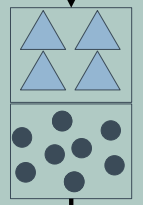

- Create QR code with Seequent Central URL
- Scan QR code using Microsoft HoloLens 2 device running the solution
- Retrieve object blobs from Seequent Central
  - Topography
  - Geological Models
  - Faults
  - Drillholes
  - Custom Meshes
  - Cross Sections
- Convert blobs into usable object data
  - Triangles
  - Vertices
  - Colours
  - Textures
- Combine data into shapes
- Render shapes over QR code

User Feedback

QR reading can be inconsistent
Requiring prepared models limits flexibility
Using public scenes threatens intellectual property

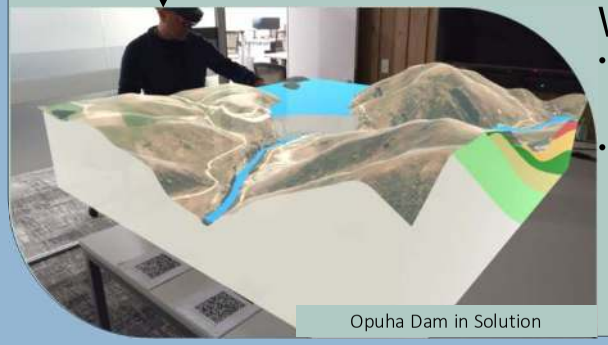

Opuha Dam in Solution

What's Next?

- More objects
  - CAD models
  - Points
- Online collaboration
  - Scan code to start session
  - Control session on Seequent Central

UC University of Canterbury — Computer Science & Software Engineering

SEEQUENT CENTRAL

Gamification in SQL Education

Goal

To **investigate** the **effects** of **competitive** and **non-competitive gamification** on student **motivation** and **performance** in learning SQL.

Background

- **Gamification** is the **incorporation** of **game-like elements** into **non-game settings.**
- It has shown to **boost performance** and **motivation**.
- Used in Education, Business and Health.

Query Competition

- **Query Competition** (QC) is a **game** based **web application.**
- It allows **students compete** by **solving queries**.
- **QC** has **challenges**, **points**, **leaderboard**, and **achievements.**

Method

- **75 students** used **QC** with **3 distinct environments**:
  - **Competitive**
  - **Gamified**
  - **Control**
- **Performance** and **motivation** were evaluated based on **retention**, **correctness**, and **speed**.

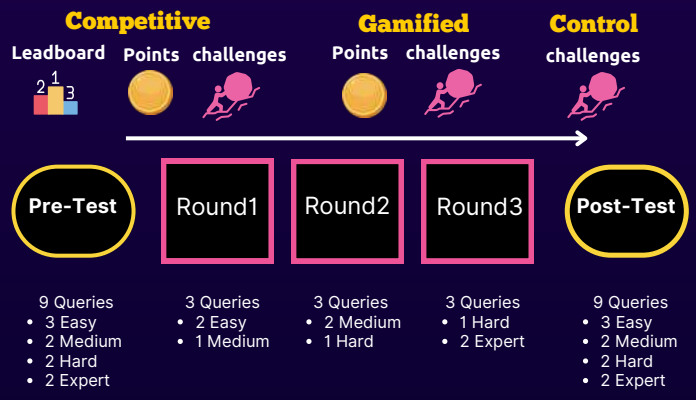

Results

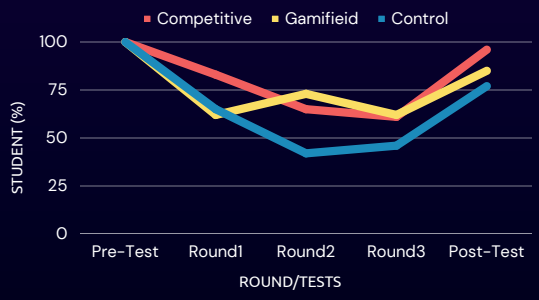

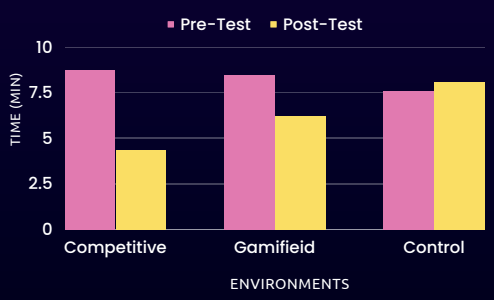

Conclusions

- **Gamified** and **Competitive groups** had **better retention.**
- **Gamified and Competitive groups had a significant improvement in answer speed.**

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Computer Science & Software Engineering

Student: **Samuel McMillan**

Supervisor: **Miguel Morales**

Mobility Data for Biosecurity

Author

Karl Moore

Supervisor

James Atlas

UC UNIVERSITY OF CANTERBURY

Computer Science & Software Engineering

Related literature

References can take up a lot of space, so cite only the key references used in the study.

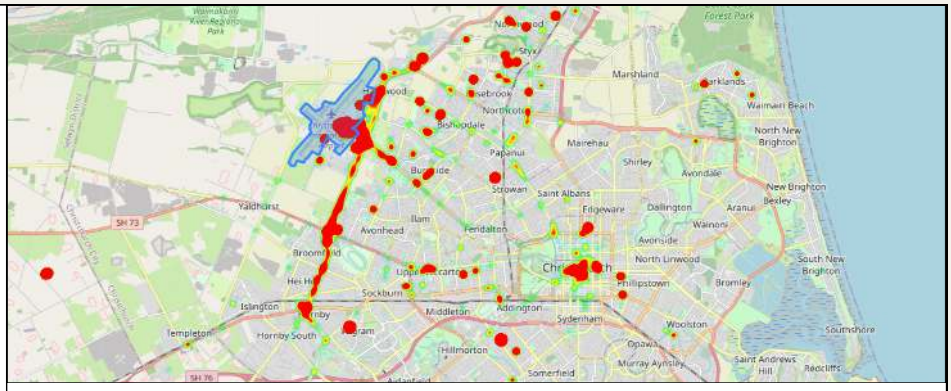

OBJECTIVE

To create an application that can efficiently and accurately display the travel patterns of those passing through bio-security risk locations. This will aid researchers in identifying potential risk hotspots for intervention.

MOTIVATION

Current border measures to exclude invasive species are imperfect, so fast identification of border penetrating threats is vital to minimise damage to the ecosystem. Threats to New Zealand's bio-security are brought from tourism and trade which create three main sources; cargo, ships, and people.

METHODOLOGY

Using an Agile development approach, a web application was built using Flask as a framework with Folium mapping libraries for spatial data display. A web application format was chosen because of the size of the data. As nine months of data for New Zealand alone is 300GB, it would create storage challenges for a desktop application.

RESULTS

The application was successfully created in a way where queries on date and coordinate ranges can be executed in an efficient manner (less than five seconds. For performance reasons, the heatmap is populated using a sample of the relevant points. Should more detailed results be required, the user is given granular control over what portion of the data is used.

Analysis

As the application is required to handle large amounts of data, performance was the major concern. The most performant data structure for the context is an R*-tree, a variation of an R-tree with a higher construction cost but faster query times for the dataset. Raw results from queries can be larger than 1GB, so sampling is necessary for timely transfers. In cases where required results are unreasonably large for transfer, a virtual desktop can be used to view the application directly on the server.

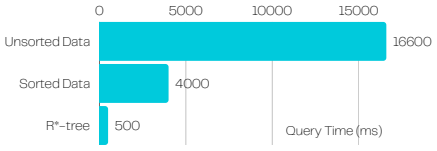

Time taken to find all individuals who visited the Christchurch Airport on a single day

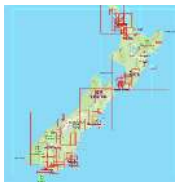

R*-tree visualisation

Conclusion

The resulting application is an efficient and accurate tool for visualizing the travel patterns of individuals passing through bio-security risk zones. This efficiency was due to the use of an R*-tree, a variation of an R-tree with a higher construction cost but faster query times for the dataset.

AUTONOMOUS DRONE FORMATIONS

Data Dissemination Protocol Project

Background:

Drones offer a unique platform for applications in areas that would be otherwise unreachable. However, sometimes a single drone is not a large enough platform for a given application. This indicates the potential usefulness of an autonomous formation of drones for use cases where a single drone is insufficient.

To successfully operate a rigid drone formation, each drone must know the positions of every other drone in the formation. This is facilitated by the Data Dissemination Protocol which transfers position, speed, direction and rotation data for every drone in the formation.

A formation of drones supports a variety of use cases including, search and rescue, biosecurity tracking and as a mobile radar array. With a single drone these use cases would be near impossible, but employing an autonomous formation not only makes them possible, but also adds benefits particularly for mobile radar applications.

Objectives:

- Successfully develop and test an autonomous drone formation improving on the existing platform.
- Demonstrate the potential of a rigid autonomous drone formation for a variety of applications.
- Develop a computationally efficient Data Dissemination Protocol as required for a successful rigid drone formation.

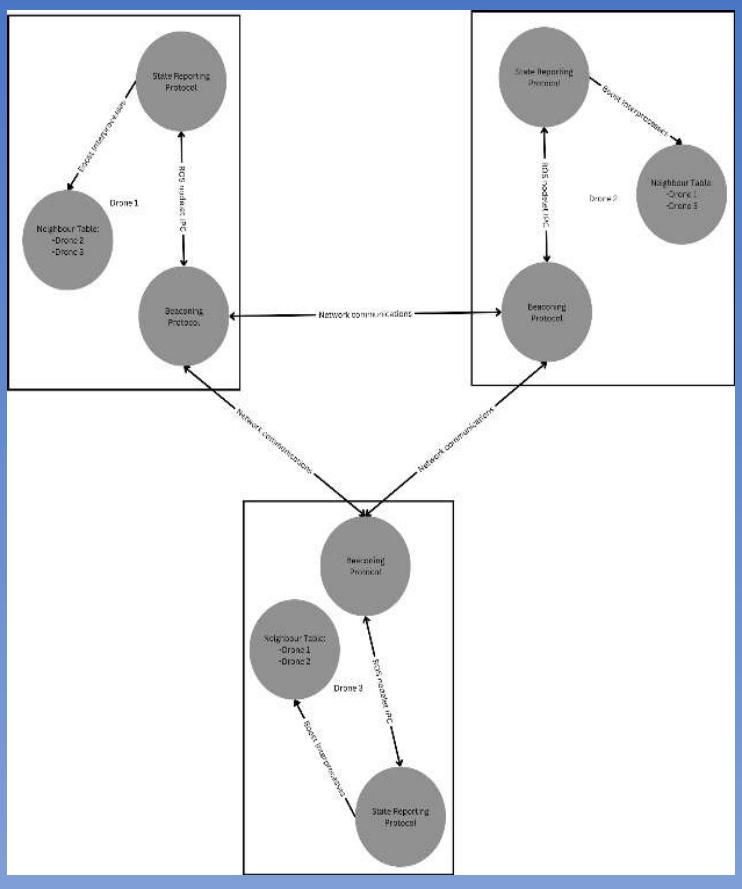

Data Dissemination Protocol:

The solution protocol provides communication between all drones in the formation and manages a neighbour table for each drone. The neighbour table stores the most recent position update received from each other drone in the formation. This allows the control loop to calculate appropriate paths for each drone to avoid in formation collisions.

The Data Dissemination protocol is made up of two concurrently running processes, the Beaconing protocol and the State Reporting protocol.

The Beaconing protocol handles transmission and reception of packets between drones, it also collects and returns data to and from the State Reporting protocol.

The State Reporting protocol interfaces with the drone control system to collect state data for transmission. It also collects received state data from other drones and populates a table of registered neighbour drones and their current state.

The Data Dissemination protocol is developed in C++ and is thread managed using ROS nodelets to ensure high throughput data flows with no inter-process copy cost. The neighbour table is stored is stored in a memory mapped file with mutual exclusion control to allow efficient, concurrent access from two processes.

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Computer Science & Software Engineering

Students: Euan Morgan, Freddy Markwell, Ethan Ng, John-Paul Lay, Josh Ellingham
Assistance: Simon Wallace-Blakely
Supervisors: Dr Andreas Willig, Dr Graeme Woodward

UC UNIVERSITY OF CANTERBURY
Computer Science & Software Engineering

Connecting Remote Families Through Storytelling

Student: Sahil Negi
Academic Supervisor: Fabian Gilson
Industry Supervisor: Tom Hughes

Problem

Many families today face the challenge of physical separation, which often leads to a lack of shared experiences and a growing emotional gap between family members. The traditional means of storytelling, which have always been a way to connect generations, struggle in this digital age to bridge the geographical divide.

Challenges:

Geographical Separation:
Modern life often scatters families across the globe, making it difficult for loved ones to share experiences in real-time.

Emotional Disconnect:
Families yearn for deeper emotional connections, and physical distance can hinder the nurturing of these bonds.

Limited Collaborative Activities: The few options available for collaborative activities across distances often lack the warmth and personal touch of traditional shared experiences.

Library Screen: Serving as a digital bookshelf, the Library screen conveniently displays all uploaded PDF books, ensuring easy access to the digital collection.

User-Centric Design: Our app is designed for users of all ages, from children to grandparents. It reimagines the way families interact across distances, merging traditional and contemporary elements. Families can now connect over shared stories and experiences, fostering stronger emotional bonds in the digital age.

Solution

Introducing a versatile mobile application, available on both iOS and Android platforms, designed to address the challenges posed by physical separation among families while fostering meaningful connections through shared stories and experiences.

Key Features:

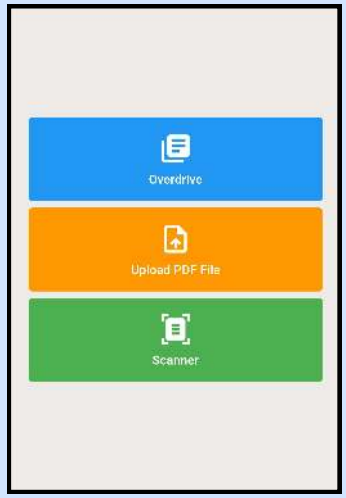

Overdrive Integration: Seamlessly integrating with Overdrive, our app lets users access and read a vast digital library, from classic literature to the latest bestsellers, enabling them to expand their literary horizons in real time.

PDF Upload:
Empowering users to curate their own collection of books, from family recipes to personal narratives, preserving treasured stories for generations. Users can directly upload PDFs to the app.

Scanner Functionality: Our innovative scanner effortlessly transforms physical pages into digital format. Scan pages from physical books, and the app converts them into accessible PDFs, anytime, anywhere.

Video Calling:
Our app facilitates real-time face-to-face interactions between friends, bridging geographical gaps and making storytelling sessions personal and engaging. Importantly, during video calls, users can seamlessly switch to other app sections, such as the Library, for a multidimensional experience where reading and connecting coexist harmoniously.

Friendship System: Building connections is at the heart of our app. Send friend requests to family members and loved ones using their email addresses. Once accepted, these connections become friends, enabling rich interactions.

GEOGRAPHIC INFORMATION SYSTEMS (GIS) ASSETS IN AUGMENTED REALITY

MOTIVATION

- **The Issue**: Most **GIS maps** are in **2D**, making them **less immersive** and **interactive**, especially for non-experts. This **limits user experience** and **data acquisition**. There's a real **need** for something **more user-friendly** and **immersive**.

- **Why It Matters**: **Transforming 2D maps** into **3D Augmented Reality (AR) interfaces** can offer **real-time location** and **position updates**, enhancing user experience significantly. It's about making **geographical information** more **accessible** and **understandable**, even in places with weak GPS signals.

OBJECTIVE

- Development of **AR Core Application** for android users that allows users to visualise their current location and position along with horizontal and vertical accuracies.

- Enhancing user interaction with **geospatial data** by allowing users to place a marker and visualise it in the AR. These markers can be used to **replace** physical markers made with **spray cans** for marking out **underground services**. This will **reduce confusion** caused by multiple markings and **eliminate environmental** concerns related to the use of paint.

- **Enhance** vertical and horizontal **accuracies** using Trimble's Global Navigation Satellite Systems (**GNSS**) with **Machine learning** on Google Streetview.

Mobile GPS & Machine learning

- **Horizontal Accuracy - 0.76 metre**
- **Vertical Accuracy - 0.52 metre**

Trimble GNSS & Machine Learning

- **Horizontal Accuracy - 0.15 metre**
- **Vertical Accuracy - 0.2 metre**

RESULTS

- Developed an **Android application** with **real-time updates** of user's current position and location.

- Provided an **AR view** for users to visualize markers placed by them.

- **Improved** vertical and horizontal **accuracies** using Trimble's GNSS with machine learning on Google Streetview.

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Trimble

STUDENT Prableen Oberoi
ACADEMIC SUPERVISOR Fabian Gilson
INDUSTRY SUPERVISOR Stuart Ralston

Natural Language Processing for Licensing Applications

Oliver Johnson, Ben Adams (Supervisor)

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

The Situation

Members of the public can object to an alcohol licence. When this happens, the District Licensing Committee (DLC) can hold a hearing where the objector(s) and applicant make their case.

From 2014 to 2022, approximately **80% of hearings result in the alcohol license being granted, despite objections**. Public objections often end up being ineffective. Is there a way to make these more effective?

The objective of this project is to research the efficacy of various natural language processing (NLP) methods applied to the decisions made by the DLC.

The end-goal is to **identify ways that the public might improve the quality of their objection** and argument in court. With this project, we try to assess if NLP techniques can be used to achieve that end-goal.

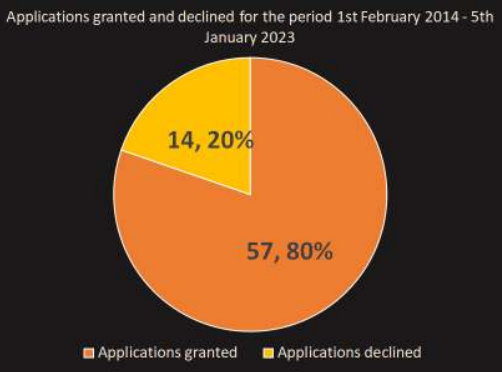

The Approaches

We analysed all the Christchurch DLC hearing decision documents from Febuary 2014 to January 2023. Later (for the topic modelling) we expanded our data to include all of Auckland's DLC hearings from January 2019 to May 2023 (lots of data!).

All of these documents are publicly available, so we downloaded them and converted them to text.

1-Word LDA Topic Modelling

We analysed the documents to see if certain topics are a strong component of the hearing, as they might have some impact on its outcome. To do so, we performed LDA topic modelling.

- Loaded data, without splitting into paragraphs
  provides larger, but fewer documents
  stripped punctuation and meaningless words
- Tokenised data
  tracked individual words and the documents they came from
- Processed tokens into bag-of-words format
  collated together into a "corpus"
- Model built and trained on corpus
  looks for 4 different topics
  3 variations: all cases, only declined cases, only granted cases

Findings:

There were not many topics in these documents and there was significant overlap between the topics. This suggests that either these hearings stay *very* on-topic, or that a lot of the nuance is lost when the hearing gets summarised in a decision.

1-Word Topic Distribution

Detailed view showing topics, their size and their overlap, as well as common words within them.

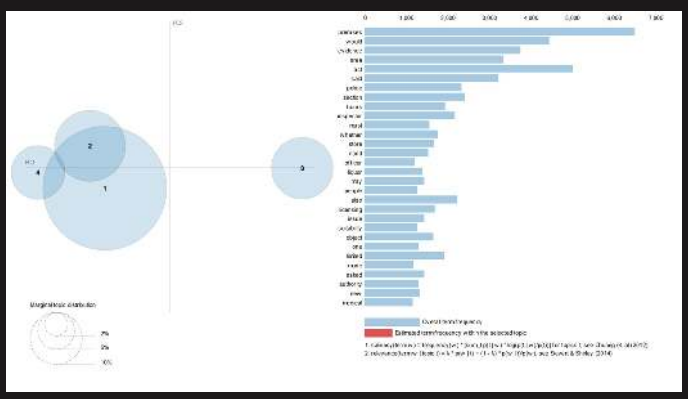

Naive Bayes

Naive Bayes is a simple machine learning algorithm. It can give us rudimentary insights, for example if the language is similar between applications that were granted and applications that were declined.

- Loaded data and split into paragraphs
  provides a higher number of data points
  prevents the algorithm from cheating by only looking at the result of the hearing.
- Partitioned the data into training and test sets - a 75-25 split
- Processed the training set into a bag-of-words sparse vector
  the vector counts each word and how many times it appears
  stripped punctuation and meaningless words
  normalised the vector using term-frequency times inverse document-frequency (tf-idf)
- Trained a classifier on vectors and evaluated
  tested classifier 50 times, using different partitions of train/test sets

Findings:

Ultimately, we found that the language used in hearing documents was incredibly similar.

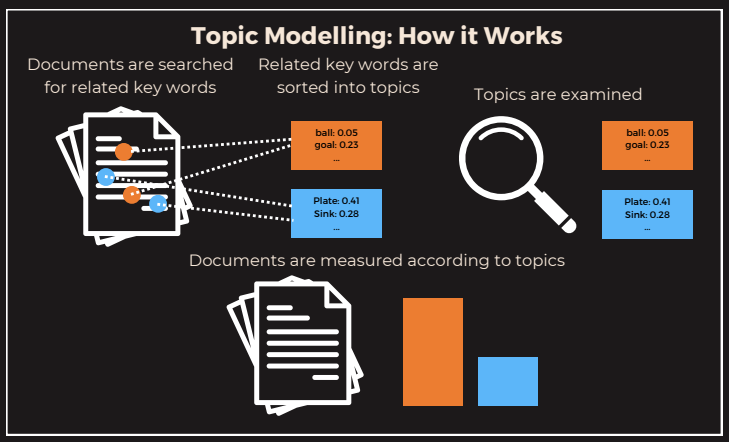

3-Word LDA Topic Modelling

Important phrases can tell us a lot about a piece of text, and the topics it discusses. To analyse these phrases, we performed some more LDA topic modelling, tweaking it to only look at phrases of 3 words.

- Loaded data, without splitting into paragraphs
  stripped punctuation and meaningless words
- Tokenised data
  tracked 3-word phrases and the documents they came from
- Processed tokens into bag-of-words format
  collated together into a "corpus"
- Model built and trained on corpus
  looks for 4 different topics
  3 variations: all cases, only declined cases, only granted cases

Findings:

Again, we found there were not many topics in these documents, confirming our findings for 1-word topic modelling.

HIP INSERTER ANGLE TRACKING

Context

Hip replacements are becoming **increasingly common**, especially for younger patients, so it is important to **improve the reliability** of the surgery. This project proposes a system that can assist surgeons during hip replacements to achieve this. The system provides the surgeon with **real-time pitch and roll angles** of a hip inserter.

Objectives

- Develop a system that calculates and displays the **pitch and roll angles** of a hip inserter from a real-time video feed.
- The system should require **no modifications** to the hip inserter.
- Improve the system so that it works accurately while the hip inserter is **partially occluded**.
- Display the pitch and roll angles **smoothly** so they can be easily read.

Angles
Pitch: 0°
Roll: 7°

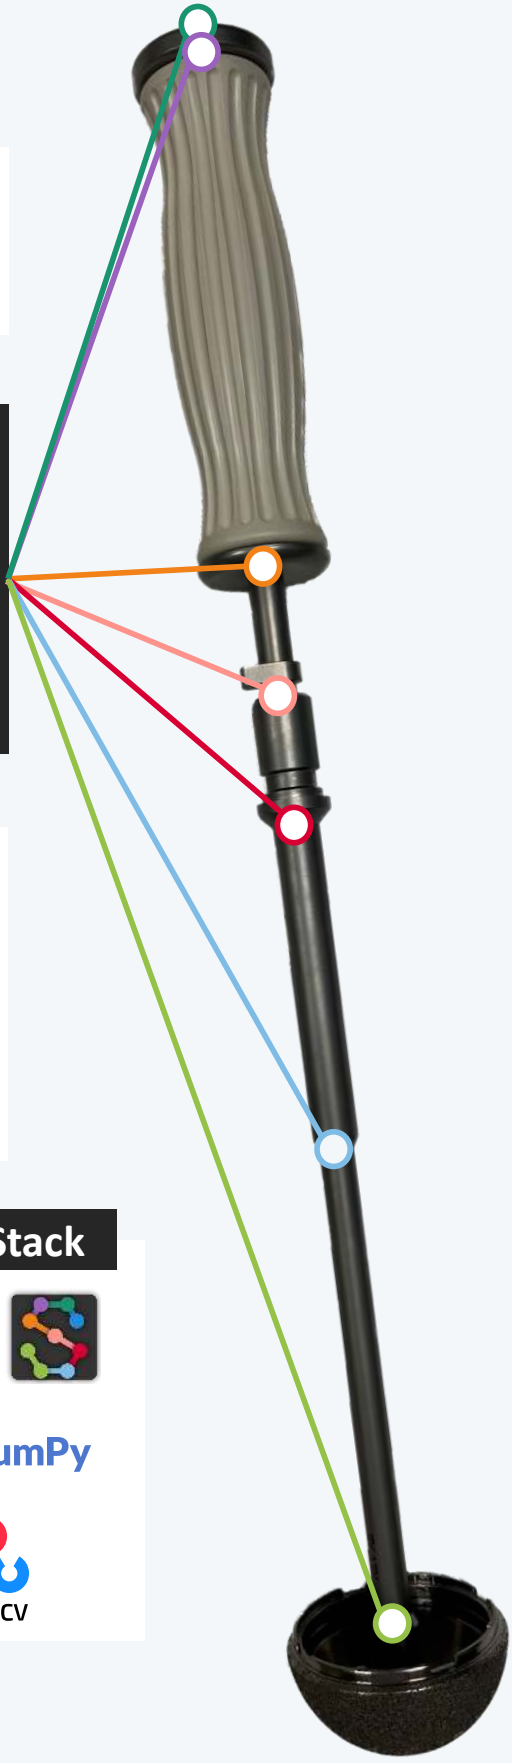

Solution

- A **ZED 2 Stereo Camera** is used capture left and right video feeds.
- A neural network trained by SLEAP is used to detect **7 chosen feature points** on the hip inserter.
- An Unscented Kalman Filter **predicts the location** of any **missing points**.
- The **2D feature points** on the left and right video feeds are converted to **3D positions**.
- A **line is fitted** through the 3D feature points and its **pitch and roll angles are displayed**.
- A Moving Average Filter **reduce jitters** when displaying angles.

Conclusion

- Similar accuracy as previous research that required adding visual markers to surgical instruments.
  - **Average pitch error: 1.0°**
  - **Average roll error: 2.1°**
- Only **2 feature points** need to be visible to display angles.
- Runs at **8-9 frames per second** with a **delay of 3 seconds** used to smooth the output.
- Current system is trained on **425 images** but needs to be trained on footage from a specific environment to work well in that environment. Future work is needed to gather **varied footage** to train a **general system**.

Tech Stack

NumPy

OpenCV

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering

Student
Dillon Pike

Supervisors
Richard Green, Casey Peat

Eyes in the sky:

A new attempt at power line detection aboard drones.

Inspecting power lines is **hazardous for humans**. UAVs present an opportunity to **make power line inspection safer**. Previous research has shown success in power line detection and guidance techniques utilising hybrid LiDAR-Vision systems and event cameras. We investigated an experimental method for **power line detection and depth estimation** that leverages **fixed-frame-rate stereo cameras**.

Method

Our method is comprised of two stages: a **power line edge detector,** and a **time aware line detector.**

The edge detector is adapted from a neural network developed in 2019 by Zhang et. al [1]. The network outputs an **edge probability map** where each pixel in the input image receives a probability of 0 to 1 to indicate the likelihood that the pixel is a power line edge pixel.

The time aware line detector is adapted from a line detector intended for event cameras, developed by Everding et. al [2]. Lines are detected by placing edge pixels into a 3D space comprised of dimensions X, Y and time. Edge pixels belonging to a line form plane structures as in *Fig. 1*.

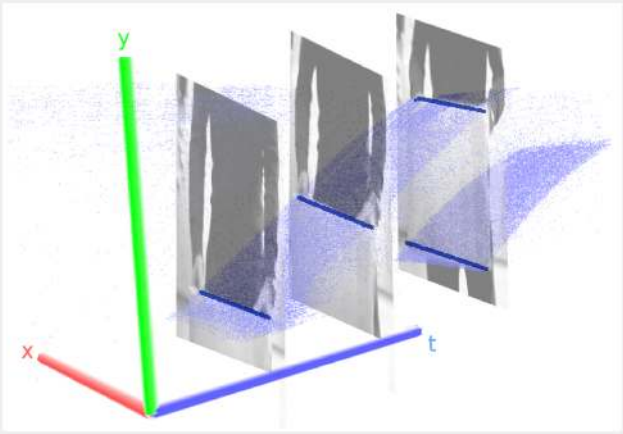

Fig 1. Example visualisation of the line detection from Everding et. al. [2]

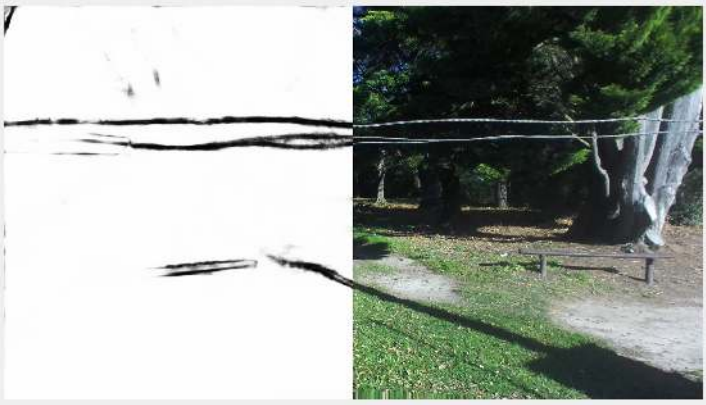

Fig 2. Sample edge detection (left) and input frame (right).

Results

Our edge detection neural network achieves competitive results compared to the pre-trained model from Zhang et. al. On sample data collected outdoors, such as *Fig. 2,* the edge detector correctly detects power lines with **high recall**, but **less than ideal precision**. When using sample data collected in a lab environment, recall remains high but precision degrades further.

The Everding et. al method for line detection is **complex, and non-trivial to implement**. Our best effort attempt at replication in Python is able to successfully identify and track power lines in scenarios where the edge detection input is precise, and the camera motion is controlled. Under most other conditions the **line detector fails to track lines.**

Key Findings

- The neural network based power line edge detection method struggles in urban environments.
- Everding et al. method for line detection does not translate easily to fixed-frame rate cameras.
- Future research using similar methods must utilise higher frame rate cameras to avoid motion challenges.

Author: Hugo Reeves

Supervisors: Sam Schofield, Richard Green

[1] H. Zhang, W. Yang, H. Yu, H. Zhang, and G. S. Xia, "Detecting power lines in uav images with convolutional features and structured constraints," Remote Sensing 2019, Vol. 11, Page 1342, vol. 11, p. 1342, 6 2019.

[2] L. Everding and J. Conradt, "Low-latency line tracking using event-based dynamic vision sensors," Frontiers in neurorobotics, vol. 12, 2018.

UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Computer Science & Software Engineering

Fusion of GNSS with Mobile SLAM

Billy Sandri

What can be improved?

- **TerraFlex** is a mobile application which allows users to **collect positional information** using a **GNSS receiver**.
- When under thick tree cover or in-between buildings the **accuracy** of the **receiver decreases**.
- TerraFlex needs **higher accuracy positional data** in a **GNSS degraded** environment.

What is the solution?

- The **camera position** of the **mobile phone** is **estimated** using Simultaneous Localisation and Mapping.
- The **camera position** is **aligned** with the **GNSS position**, to be in the same **coordinate frame**.
- The user is provided with a '**Fused**' **position** where an engine switches to the **camera position** when **GNSS** accuracy **degrades**.

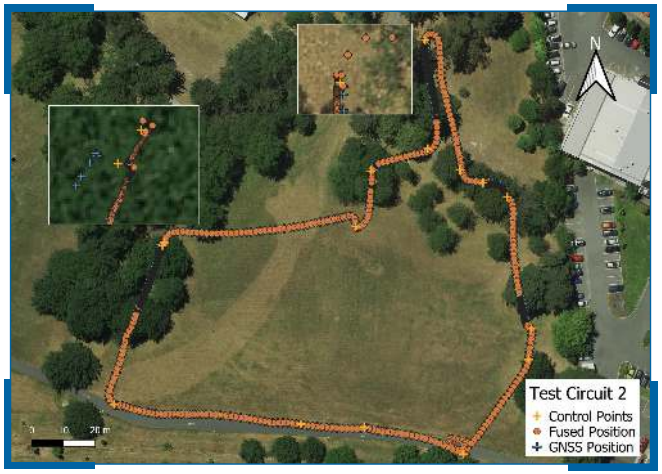

What is the outcome?

- There is **evidence** that the '**Fused**' **position** can provide **higher accuracy positional data** to the user in a **GNSS degraded** environment.

Supervisor
Fabian Gilson

Industry Partner
Stuart Ralston

Trimble

TerraFlex

UC UNIVERSITY OF CANTERBURY
Te Whare Wānanga o Waitaha
CHRISTCHURCH NEW ZEALAND

Computer Science & Software Engineering

Aerial Precision: LiDAR Conductor Detection

Motivation

- Unison owns and maintains over 9000km of power lines
- Manual conductor inspection is time consuming and expensive
- Drones offer a promising alternative to manual inspection
- Unison found human-operated drones hard to position optimally
- Automated drones are more precise than manual control

Method

- Filter out data far away from where a line was detected in the last detection
- Filter out noisy data using a k-nearest neighbours
- Find lines in the data using a customised RANSAC algorithm
- Find groups of parallel lines and calculate which is the best
- Identify which line from the group is the one to track
- Output a vector for a drone to travel along to stay above the line

Results

- All three conductors accurately detected in 100% of frames
- Point classification **accuracy of 99.98%** ensures reliable conductor detection
- **Precision of 100%** indicates no points falsely classified as conductor points
- **Recall of 91.57%** shows the algorithm's ability to classify conductor points
- **F1 score of 95.66%** signifies an excellent balance between precision and recall

Future Work

- Improve algorithm for intersecting and overlapping lines
- Investigate combined LiDAR and camera systems to improve conductor detection
- Drone integration and real-world testing

UC University of Canterbury Te Whare Wānanga o Waitaha Christchurch New Zealand

Computer Science & Software Engineering

Student
Bede Skinner-Vennell

Supervisors
Richard Green
Sam Schofield

Sponsor
unison
Alex Castellanos
Ben van Vliet

MOBILITY DATA VISUALISATION

A TOOL FOR DETECTING SUSPICIOUS GEOSPATIAL PATTERNS

AMY SLOANE

INTRODUCTION

In the realm of geospatial research, the visualization of mobility data stands as a potent tool, offering insights into intricate human movement patterns. While extensively applied in transportation, its potential within security and policing remains an underexplored frontier.

This research initiative introduces a compelling proof of concept - an innovative web-based application designed to flag buildings experiencing an unusually high influx of individuals. Through meticulous research, development, and evaluation, this research seeks to pave the way for a more secure and data-driven future.

MOTIVATION

"How can we utilise Mobility Data - Visualisation in the unexplored domain, ***Security****?"*

EXAMPLE MOBILITY DATA

tsv file format

	Hashed_ID	lat	lon	Point_ID	date	time
749638	abc12345	-43.533143	172.612229	1633481940	2021-10-06	13:59:00
749661	abc12345	-43.535798	172.612239	1633488072	2021-10-06	15:41:12

PROJECT METHOD

- Research
- Plan
- Ethics Application
- Build
- Test
- Refine
- Evaluate

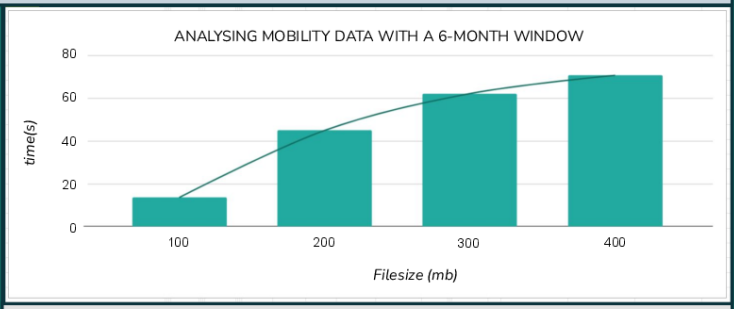

TECHNOLOGY STACK

python · Folium · GeoPandas · django · dask

RESULT

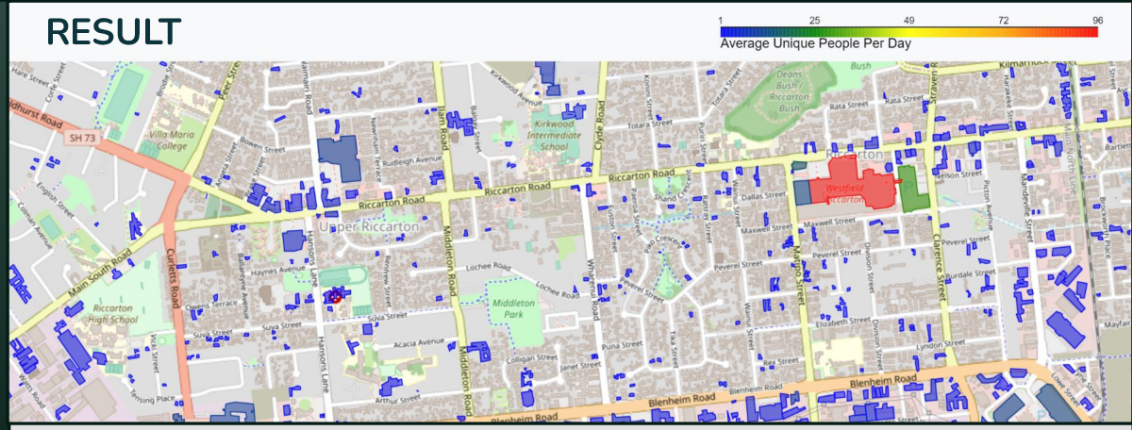

KEY QUALITY ATTRIBUTES

- Maintainability
- Scalability
- Accuracy
- Security
- Performance
- Usability

CORE ALGORITHM

Find suspicious/unusual geospatial patterns by highlighting properties with a large volume of unique persons within a time period

Step 1: Refine Mobility Data by Date and Time
Step 2: Merge Mobility Data with Building Data (LINZ NZ Building Outlines)
Step 3: Remove all buildings that don't intersect with mobility data, and vice versa
Step 4: Merge matching anonymous IDs together within a building
Step 5: Calculate Average Number Unique Persons Per Day
Step 6: Display Result on a Map using relevant colour scaling

EVALUATION

Performance and Scalability

- Performance tests were conducted on a sample dataset, measuring just under 400MB.
- Results indicate a linear relationship between performance and the size of the input mobility data.
- Displays promise in executing its algorithm efficiently with medium-sized mobility data files, but it's not yet capable of handling larger sizes (100GB+).

Limitations

- Demonstrated a general linear performance trend, but occasional wide (30 second) variances were observed
- Further testing with larger files required to make accurate performance predictions using the complete dataset

Conclusions

- Achieved a successful outcome in generating a map that highlighting buildings with unusually high person-traffic.
- Indicates potential for future projects in detecting unusual or suspicious patterns. These advancements have the potential to provide valuable support within the security sector.

ACKNOWLEDGEMENTS

VANESSA BASTOS
SUPERVISOR

JAMES ATLAS
SUPERVISOR

UC UNIVERSITY OF CANTERBURY

Scott Base Seal Census Using Machine Learning

Jacques Terblanche

Why Count Seals?

Counting Weddell seals at Scott Base, Antarctica, isn't just science—it's a necessity with far-reaching implications. In Antarctica's extreme conditions, manual tracking is impractical. Machine learning offers a transformative solution, enabling rapid, accurate, and non-invasive assessments, a crucial advantage during Scott Base's extensive reconstruction.

The driving force is mitigating construction noise's potential impact on seals. Even distant sounds stress these creatures. Continuous, precise monitoring, aided by machine learning, ensures seal welfare and aligns with Antarctica New Zealand's conservation commitment.

Beyond construction, historical data informs broader seal behavior effects. Yet, harsh conditions make manual data collection impractical. An automated system for robustly detecting seals in images is essential.

Two key convolutional neural networks (CNNs) come into play. RetinaNet identifies seals in images from past summer seasons (2018-19 to 2022-23), while a Residual Network (ResNet) detects snowstorms, aiding in false low seal count correction. The project's primary goal: enhancing these CNNs' performance in seal population monitoring.

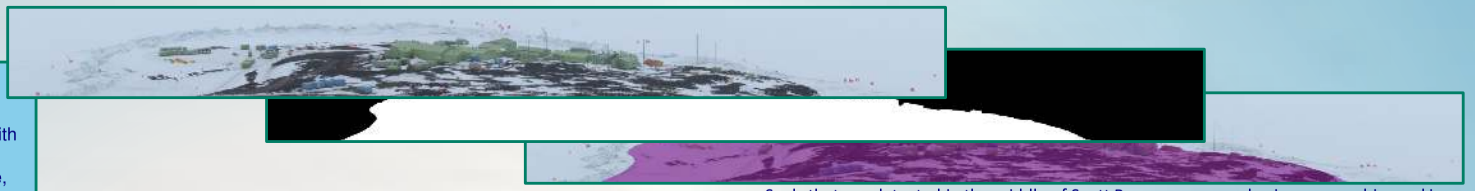

Seals that are detected in the middle of Scott Base are removed using geographic masking

Spotting Seals with RetinaNet

Our seal detection relies on the robust RetinaNet CNN, purpose-built for identifying small, densely packed objects like seals. RetinaNet employs two key innovations: Feature Pyramid Networks and Focal Loss.

Feature Pyramid Networks: Ensures precise detection of small, distant seals by harmonizing detailed low-resolution and high-resolution features. This innovative approach allows RetinaNet to adapt to seals at varying distances within our challenging Antarctic environment.

Focal Loss: Bolsters the model's performance in challenging conditions, prioritizing difficult detections. It ensures seals receive attention amid complex backgrounds.

In essence, RetinaNet integrates these techniques, creating an accurate and efficient seal detection system. The network's architecture synergizes a bottom-up pathway for scale, top-down pathway and lateral connections for context, a classification subnetwork for seal identification, and a regression subnetwork for precise location.

This approach allows RetinaNet to identify seals with exceptional sensitivity, capturing anything resembling a seal, be it actual seals, rocks, or cracks in the ice. Detections are filtered to find actual seals.

Filtering Out the Noise – Eliminating False Positives

To enhance the accuracy of our seal detection, we employ two key filtering techniques.

Geographic Masking: Leveraging geographic data to differentiate true seal detections from false positives. We mask out Scott Base and surroundings, excluding land detections as seals primarily inhabit sea ice near ice cracks. This binary mask effectively filters out any land-based detections.

Temporal Comparison: Further filtering involves analyzing detections in adjacent images. Missing counterparts in neighboring images flag potential false positives. This temporal comparison not only boosts result precision but also eliminates false detections in outlier images affected by sun glare, which may highlight ice cracks, resembling seals.

These techniques, combined with RetinaNet's exceptional sensitivity, ensure accurate seal population identification and tracking.

Tiny Seals, Vast Expanse: Region of Interest

The camera used to capture our dataset is positioned far above Scott Base to avoid construction activity while also covering the whole seal colony. In an image the size of this poster, our region of interest would be this area. Can you spot the seals in this image?

Seal locations verified using detections from neighboring images

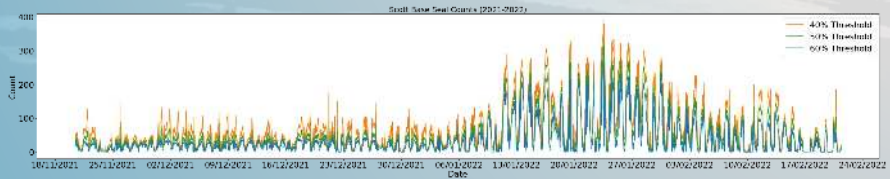

Visualizing Seal Counts: This line graph depicts seal counts during the 2021-22 season at confidence thresholds of 40%, 50%, and 60%. It highlights our model's proficiency in tracking seal populations under varying confidence levels, demonstrating the effectiveness of our methodology.

Results – Measuring Progress

Our seal detection enhancements have yielded promising results. Comparing the 2022 and 2023 models reveals significant progress. The 2023 model, with masking, notably improves precision, accuracy, recall, and F_1 score over the 2022 model. Introducing filtering further boosts accuracy while slightly impacting other metrics, crucial for reducing false positives beyond the mask.

F_1 score, which balances precision and recall, provides an overall model performance assessment. The 2022 model achieved an 87.1% F_1 score, while the 2023 model with masking and filtering reached 91.6%, signifying substantial precision improvement without compromising recall.

These results affirm our success in enhancing seal detection precision and tracking at Scott Base. Beyond our immediate goals, this data supports Antarctica New Zealand's broader conservation efforts, aiding seal behavior study, stress factor identification, and ecosystem preservation.

Model	F_1 Score	Precision	Recall	Accuracy
2022 Model	87.1%	**98.5%**	78.1%	77.2%
2023 Model	79.7%	70.3%	**92.0%**	66.2%
2023 with masking	**91.8%**	91.5%	**92.0%**	**84.8%**
2023 with masking and filtering	91.6%	92.3%	90.8%	84.5%

UC University of Canterbury
Computer Science & Software Engineering

Sponsored by Antarctica New Zealand

Supervised by Richard Green
Advised by Oliver Batchelor

By Jacques Terblanche
jacquest258@gmail.com

UC UNIVERSITY OF CANTERBURY
Computer Science & Software Engineering

Contributor: Jamie Thomas
Supervisor: Ben Adams
Sponsor: Foam

FOAM Proof-of-Location Services at The University of Canterbury. How Does it Compare?

Background

Regular location services can be spoofed and are unverifiable given certain exploits that can be used. A company based in Brooklyn called Foam has created a decentralised proof-of-location system that can give GPS accuracy while being less susceptible to spoofing. This project involves **implementing Foam's technology** at the University of Canterbury and **evaluating its performance** compared to other standards of location services such as GPS.

"*A proof of location is a digital certificate that attests someone's presence at a certain geographic location, at a certain time.*"[1]

1. Brambilla, Giacomo & Amoretti, Michele & Zanichelli, Francesco. (2016). Using Block Chain for Peer-to-Peer Proof-of-Location.

Locations

Four roof locations were selected to install the Zone Anchors on. They were

- West,
- Rahua,
- Meremere and
- Puaka - James Hight (the Central Library).

These were selected as they were high enough to have **line of site for the antennas** needed for the nodes. These buildings also had the best roof accessibility, which helped us in surveying and installing the nodes.

Building and Testing

Four Zone Anchors (nodes) need to be built and set up around the campus to create a Trust Zone that can detect mobile nodes in an area. The materials to build these were sent from Foam and build at UC. These nodes took several weeks to build and took the most time to get working.

Debugging:

The power supplies sent with the nodes **couldn't charge the batteries** at the same time as powering the nodes. This was fixed by getting **new power supplies** that used adapters from US plugs to NZ plugs. This issue was hard to debug due to the inconsistency of the power supplies.

Test Bed

Proof of Location was a web application designed to allow testers to run experiments on the FOAM protocol's location services and compare it to other location services.

This application was made using React and used MapBox as its integrated map service.

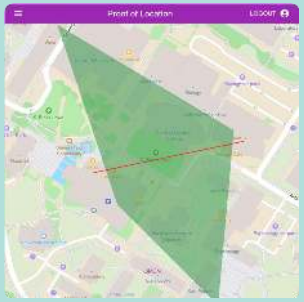

There was also an **android mobile app** designed and built for this purpose. It uses Jetpack Compose and MapBox and is a simple application for recording experiments easily.

Using Mobile Location Data for Climate Friendly Transport Interventions

Problem

- Road transport is responsible for **15%** of New Zealand's greenhouse gas emissions, and needs to be reduced to meet domestic and international targets.
- To do this, **car usage needs to be reduced** by making alternative travel viable.
- This requires good data, but current collection methods are slow and expensive. Data is collected at only one place and over a short period of time.

Objective

- Harness **mobile phone location data** to give transport researchers a quick, easy, no-technical-knowledge-required way of assessing the impact of transport interventions.

Solution

A **web app** was developed. Transport researchers can use it to make complex queries of mobile phone location datasets, with no programming knowledge required.

Tech stack

dask Flask Parquet

Steps

1: Select a map area.

2: Set query parameters.

3: View and save results.

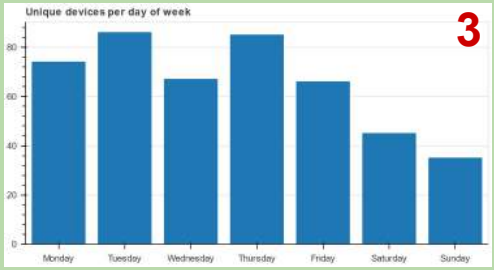

Results

The web app can find the number of unique devices in an area during a period of time. Results can be graphed over different time periods, such as days of the week. Traffic can thus be compared pre- and post-intervention to determine effectiveness.

In terms of performance, the web app can perform a complex query on **275 GB** of data in **under 45 minutes** running on **just a laptop** - and it only needs 51 GB of storage space, data included.

Advantages Over Status Quo

Comprehensive. The dataset covers the entirety of New Zealand.

Cheap. The data is buy once, analyse forever.

Fast. Analysis takes just minutes; no wait to collect data "in the field".

Easy. The web interface requires no programming knowledge to use.

UC University of Canterbury Te Whare Wānanga o Waitaha Christchurch New Zealand

Computer Science & Software Engineering

Student: Jonathan Tomlinson
Supervisors: James Atlas, Vanessa Bastos
Associated researcher: Simon Kingham

Social sustainability in the NZ software industry

By Saskia van der Peet | Supervised by Miguel Morales-Trujillo and Ismael Caballero Muñoz-Reja

context

Software has become such an integrated part of our lives, so it is important that it is sustainable.
Sustainability in software has five dimensions: economic, environmental, individual, social and technical. The **social dimension focuses on software's impact on communities.**

motivation

- **NZ has committed to achieving the UN 17 Sustainable Development Goals (SDGs) by 2030.**
- **12 of the goals are impacted by social sustainability.**
- **Until now there has been no current understanding of social sustainability in NZ software industry.**

objectives

To discover the:

1. **Factors of social sustainability** in software
2. **Frameworks for social sustainability** in software?
3. Social sustainability **practices in NZ software industry?**
4. **Challenges** of applying social sustainability practices in the NZ software industry?

method

1 literature review
- keyword search:
  - sustain*
  - software
  - engineering
  - development
  - social*
  - holisitic*

2 design survey & interview
- sections:
  - demographics
  - UN SDGs
  - social responsibility
  - organisation processes
  - challenges

3 conduct survey and interviews
- with NZ software professionals

practices of social sustainability

- **The product and organisation mission is sustainability.**
- **Organisations obtain sustainability certifications (e.g. BCorp)**
- **Employees get paid to volunteer**

challenges of social sustainability

A lack of awareness and understanding of the impact of work.

Business interests prioritised over sustainability, i.e. money or customer requirements not aligning with sustainability.

Four most considered SDGs during software development life cycle.

conclusions

NZ software organisations are primarily contributing to social sustainability through various **financial initiatives**. A **few organisations consider the impact of their product on the community**, but its extent is **constrained by money**.

Furthermore, software professionals **lack awareness, understanding of impact, and education** to practice social sustainability throughout the software development life cycle.

These challenges are not dissimilar to sustainability in fashion, food, and other industries. Consumers choose cost over **sustainability** as it **does not have a visible, direct impact on them**, unlike paying for something.

> "The **important companies** are the ones that **are doing social sustainability** as part of their reason for being."
>
> - Professor at a New Zealand polytechnic

Scan me for more in depth details!

https://marigoldsunsetdaisies.notion.site/

UC UNIVERSITY OF CANTERBURY
Computer Science & Software Engineering

Laboratory Automation at ESR

By: *Moses Wescombe*
Supervised by: *Fabian Gilson*
Partnered with: Michael Lechermann, Levi Bourke (ESR)

Problem

The Institute of Environmental Science and Research (ESR) operates a laboratory for measuring the gamma radiation of samples. Currently, the samples are **manually placed** in detectors which can take hours to complete. A **limited number of detectors** and operation hours result in the laboratory's **restricted throughput**.

Solution

Program a **software stack** alongside a **robotic arm** to manage the registration, movement and processing of samples.

- The robotic arm has four detectors within reach
- Samples are registered from a **browser control panel**
- Samples are **queued for processing** and the operation is managed by the arm
- Samples are identified using a combination of a **depth camera** and **fiducial markers**.

Outcomes

Currently, the system is able to **correctly identify** and **move samples** according to a queue. The project is currently still in progress and more work is needed to fully integrate the detectors.

Automating radioactive sample **measurement**, potentially **doubling throughput** while reducing human input.

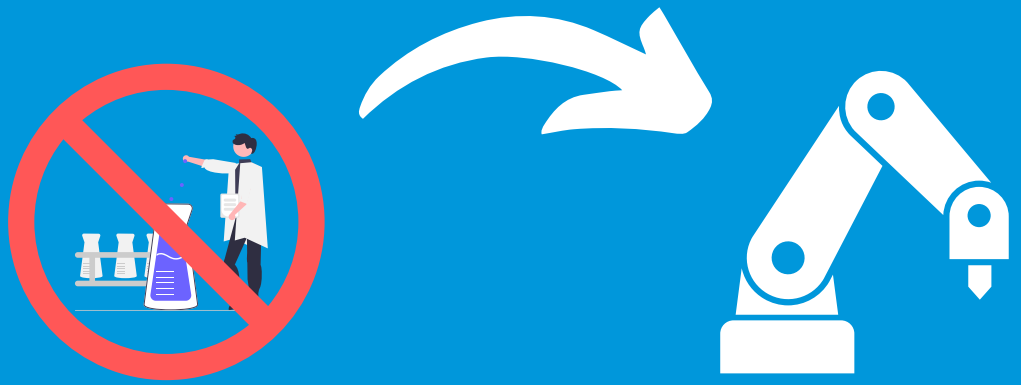

UC UNIVERSITY OF CANTERBURY Te Whare Wānanga o Waitaha CHRISTCHURCH NEW ZEALAND
Computer Science & Software Engineering
ESR Science for Communities He Pūtaiao, He Tāngata

AI TO DIAGNOSE DIABETIC EYE HEALTH

MOTIVATION: People with diabetes are at risk of losing their eyesight, but this can be prevented by prompt screenings.

Diabetic retinopathy (DR) ranks among the leading causes of preventable blindness worldwide. With diabetes on the rise, prompt access to eye screenings means eye damage caused by diabetes can be detected, and treatment can be initiated to prevent vision loss. New Zealand's current manual screening process is slow and relies on professionals for image analysis, causing bottlenecks that hinder expanding patient care.

BY 2040, OVER

600 million

PEOPLE WITH DIABETES WILL REQUIRE REGULAR EYE CHECKS

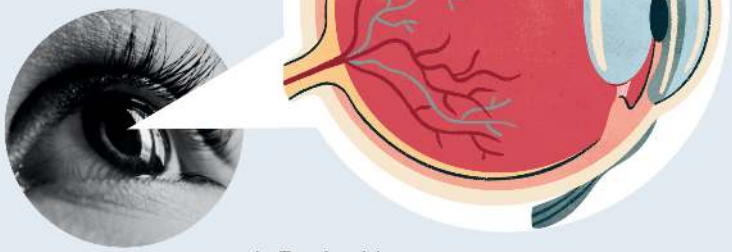

OBJECTIVE: The aim is to use machine learning to automate the DR screening process for NZ.

So, if early detection can help prevent blindness, let's improve the system so more people can access prompt screenings.

Automating screenings 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 machine learning model tailored for New Zealand.

To address this, a machine 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.

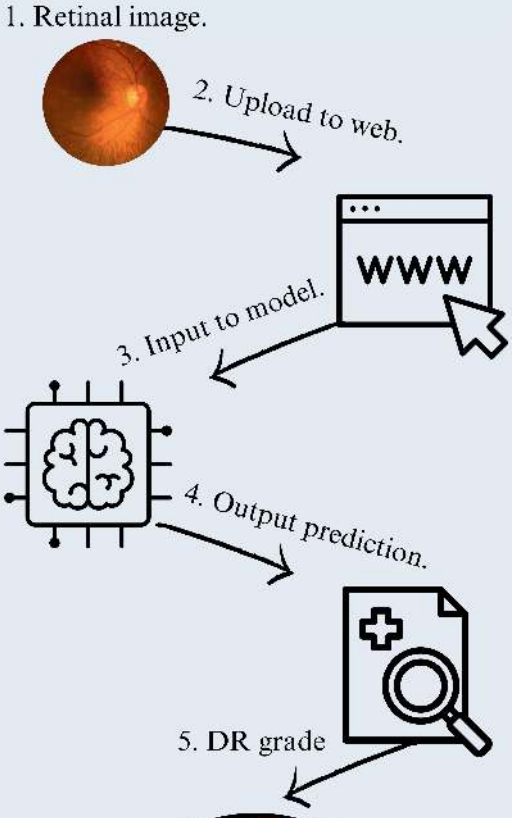

OUTCOME: <95% accuracy in predicting DR grades.

Machine Learning Model Summary

- Type: Supervised Classification.
- Method: Transfer learning with ResNet50 as the base model.
- Input: Retinal images from Auckland University.
- Output: DR grade on a scale from 0 to 5.
- Tech Stack: TensorFlow, Python, LabelStudio, Jupyter Notebook.

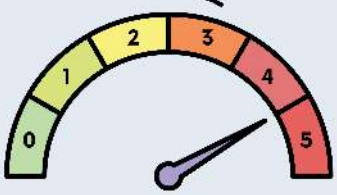

DR grade scale recommended by the Ministry of Health NZ.

0 - Healthy
1 - Minimal DR
2 - Mild DR
3 - Moderate DR
4 - Severe DR
5 - Proliferative DR

Student
Yiyang (Jessie) Yu

Supervisor
Andrew Bainbridge-Smith

UC UNIVERSITY OF CANTERBURY

Computer Science & Software Engineering

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