MTU

Ollscoil Teicneolaíochta na Mumhan

Munster Technological University

Computer Science Final Year Project

2025/26 Poster Selection

Nexus Centre,

MTU,

Bishopstown, Cork

6th May 2026

AI Vulnerability & Defence in Image Classification

Analysis of Adversarial Attacks and Defences in Image Classification

Isihaq Abass, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Deep learning models, particularly convolutional neural networks (CNNs), are widely used for image classification tasks but are known to be vulnerable to adversarial attacks. These attacks introduce small, imperceptible perturbations that can cause incorrect predictions. This project investigates gradient-based adversarial attacks (FGSM and PGD) and evaluates defence strategies using the CIFAR-10 dataset. The aim is to assess how these attacks affect model performance and to analyse how adversarial training influences model robustness within the system.

Objectives

- Train a baseline CNN on CIFAR-10
- Implement FGSM and PGD attacks
- Apply adversarial training as defence
- Evaluate performance on clean vs adversarial data
- Analyse robustness accuracy trade-off

Research Problem

"To design and implement an educational framework demonstrating how gradient-based adversarial attacks affect CNN performance and how adversarial training influences model robustness, including the trade-off between robustness and clean accuracy."

Methodology

CNN Model

- A convolutional neural network (CNN) was trained on the CIFAR-10 dataset
- Consists of convolutional + pooling layers
- Achieves high accuracy on clean CIFAR-10 test data

Adversarial Attacks

- **FGSM**: Single-step gradient-based attack
- **PGD**: Iterative, stronger attack
- Both attacks aim to maximise prediction error

Adversarial Training (Defence)

- Model retrained using adversarial examples
- Improves robustness against attacks
- Reduces performance degradation under attack

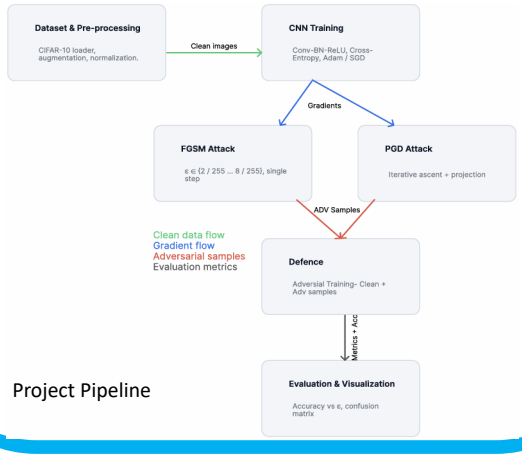

Project Pipeline

Evaluation Results

Adversarial Attack and Defence Example:

FGSM attack causes misclassification, while adversarial training restores correct prediction

Results:

Adversarial attacks reduce baseline accuracy, while adversarial training improves robustness, especially against FGSM and PGD.

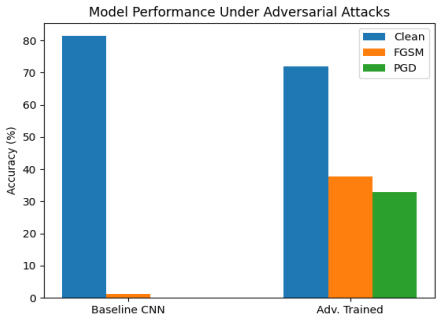

Key Insights:

- PGD attack is significantly stronger than FGSM
- Adversarial training improves robustness
- Small perturbations exploit vulnerabilities in decision boundaries
- Defence introduces a trade-off between robustness and clean accuracy
- The baseline model exhibits near-complete failure under strong adversarial attack (PGD), highlighting the vulnerability of standard CNNs.

Conclusions

- Adversarial attacks significantly reduce CNN classification accuracy, even with imperceptible perturbations
- PGD attack is more effective than FGSM, causing greater degradation in model performance
- Adversarial training improves robustness by enabling correct classification of perturbed inputs
- Trade-off exists between robustness and clean accuracy due to defensive training
- CNN models are highly vulnerable, but defence strategies significantly improve overall resilience

References

1. Goodfellow, I. et al., Explaining and Harnessing Adversarial Examples, 2014.
2. Madry, A. et al., Towards Deep Learning Models Resistant to Adversarial Attacks, 2018.

Acknowledgments

I would like to acknowledge the support and guidance of Dr Oonagh O Brien, Dr Christian Beder and the resources provided by Munster Technological University throughout this project.

AutoSec

Laptop-Only Automotive Cybersecurity Testbed

Luté Luka Abrahams, BSc Honours in Software Development, MTU Cork

Scan for Website

Attack ✦ **Alert** ✦ **Containment**

Introduction

- Modern vehicles expose cyber attack paths via gateways, diagnostics, and connectivity.
- CAN bus has no sender authentication and limited built-in security controls.
- Commercial vehicle IDS solutions are often hardware-bound and costly for students testing.

Threat Model

- Asset: In-vehicle CAN traffic and safety-relevant
- Attacker capability: Inject frames on the bus (flood/ spoof/replay).
- Trust boundary: Traffic entering the vehicle network is treated as untrusted.
- Goal: Detect and contain suspicious traffic before it reaches the vehicle.

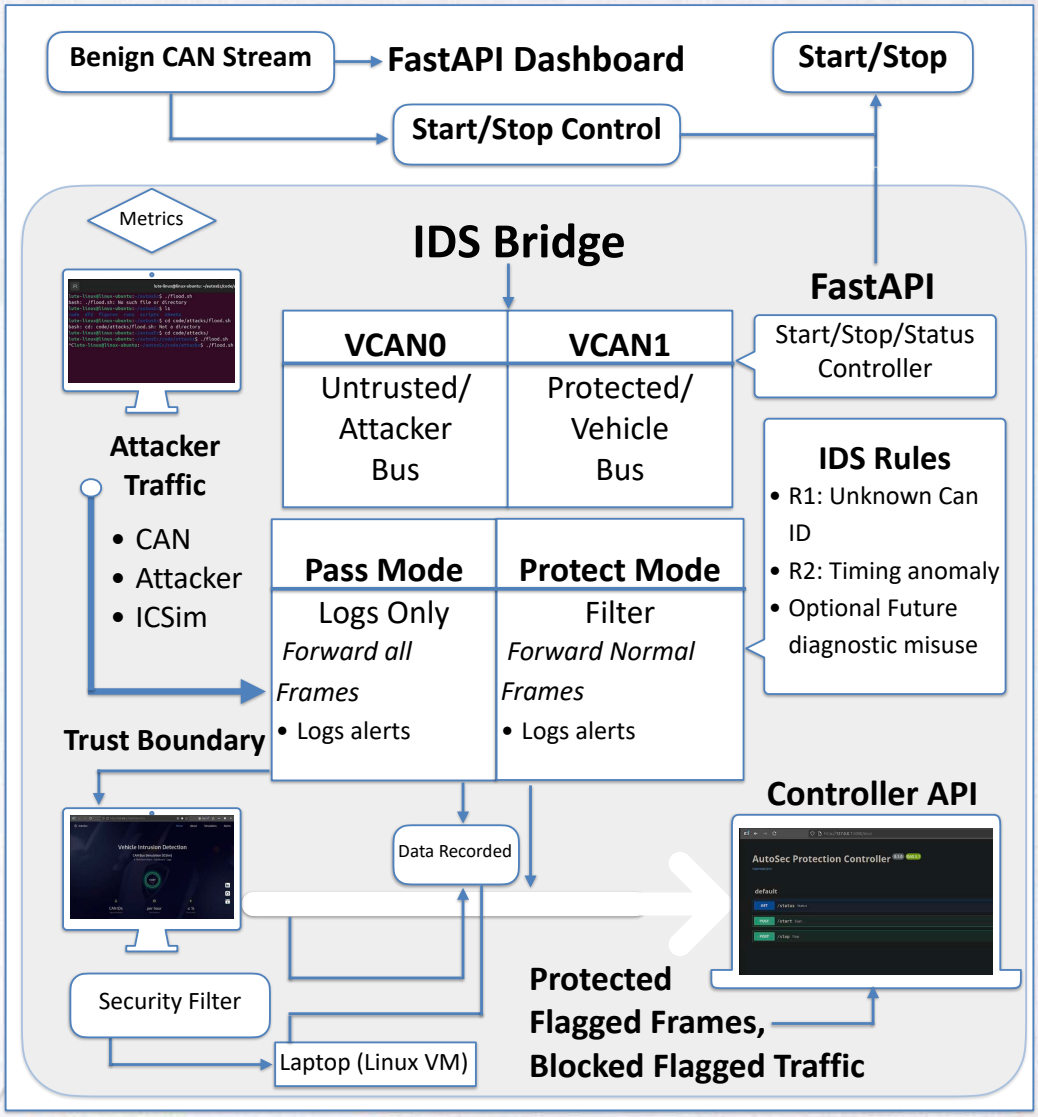

Key Results

- No Protection
  - Run ICSim (vehicle) on vcan0.
  - Inject an attack (eg, flooding) onto vcan0.
  - Result: traffic spikes and the simulated vehicle UI.

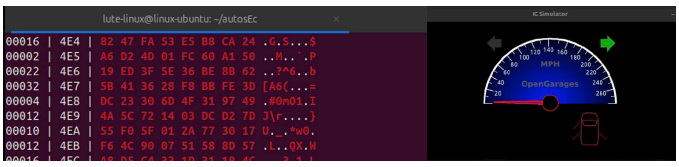

Demo Flow

- Protection Enabled
  - Run ICSim (vehicle) on vcan1.
  - Attacker injects traffic onto vcan1
  - IDS Bridge filters vcan0 → vcan1 remains clear

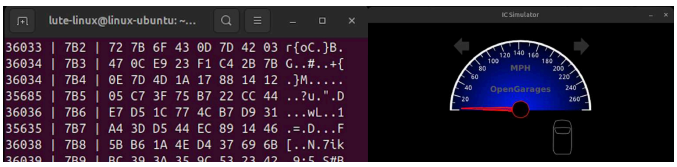

No Protection ⟶ **Protection Enabled**

References

- ISO, ISO/SAE 21434: Road Vehicles - Cybersecurity Engineering, 2021
- UNECE, UN Regulation No. 155 Cyber Security and Cyber Security, 2021.
- Miller, C., and Valasek, C., Remote Exploitation of a Passenger, 2015

Acknowledgement

- I gratefully acknowledge Dr. Paul Davin & my brother for guidance and assistance with idea

Extended Reality Notes

Aether Allen, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Forgetting is a part of everyday life, and it is a necessary function of the brain to allow for predictions of the future [1]. However, if a person forgets important tasks, it can be disruptive to that person, or whoever or whatever requires that task to be completed. The challenges for memory are distraction and failure to start tasks whether because of apathy or inability to remember them [2].
This project aims to solve these challenges by using augmented reality features, such as spacial anchors, as a base for a notes and reminders application. This application will allow a user to create notes for everyday tasks with all the necessary details including location which is intended to overcome the challenge of distraction, and time and date notifications to help overcome the challenge of forgetting.

Application Features

- **Augmented Reality:** Immerses user with notes without a mess
- **Note Creation:** User can place notes
- **Note Deletion:** User can delete a note
- **Note Editing:** User can edit details
- **Time Notifications:** For tasks done at a specific time
- **Date Notifications:** For tasks done at a specific date
- **Location Notifications:** To help overcome the challenge of distraction
- **Priority Levels:** A visual display to help user decide which tasks to complete first

Evaluation Process

Evaluation will check how the project works and how it effectively solves the problem.
Technical: Manual and code-based tests to check that features work as intended.
Task completion: The average number of tasks testers can complete with and without the app. This gauges the effectiveness of the app to help with memory.
Survey: Testers will be surveyed to see if the app reduces or increases stress levels to see if the task is solving a problem without creating a new problem.

Application

Location Based Notification

Notification Off

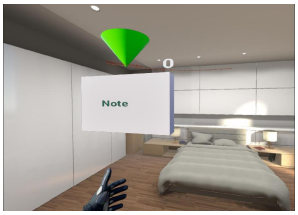

Notification On

Editing Note Details

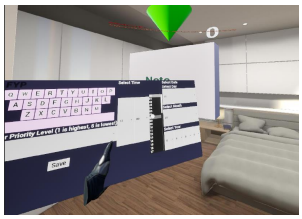

Text, priority level, and time and date notifications

Note containing information set by the menu

Technologies

Meta Quest

UNREAL ENGINE

OpenXR™

References

1. Gravitz, L., "The forgotten part of memory", Nature, 2019
2. Lehman, S., Graves, J., Mcaleer, C., Giovannetti, T., Tan, C.C., "Human interface and the management of information. Information in applications and services", Mobile Augmented Reality Game to Encourage Hydration in the Elderly ,6 2018

Acknowledgments

The author would like to acknowledge
Larkin Cunningham for supervision over the project

Edge-AI Fruit Detection

Comparison with an Industrial Vision System

Maham Aziz
BSc Honours in Software Development
Department of Computer Science
MTU Cork, May 2026

Introduction

This project investigates real-time fruit detection and classification using Edge AI. A fruit image dataset was used to train and compare lightweight deep learning models across nine classes, with testing carried out using different dataset sizes to evaluate accuracy and efficiency. The best performing models were then deployed on edge hardware and investigated alongside an industrial vision setup to assess suitability for real-world automated inspection tasks.

Problem Statement

Automated inspection systems are often built using expensive industrial vision equipment which can limit accessibility for smaller scale applications. This project investigates whether a lower-cost Edge-AI based object detection system can be trained with limited data and still deliver effective real time fruit detection and classification.

Methodology

A fruit image dataset was prepared for object detection across nine classes and used to train lightweight deep learning models under comparable conditions. Additional experiments were carried out on reduced dataset splits to examine the effects of limited training data on detection performance. The models were then evaluated using standard metrics, deployed on edge hardware and compared with an industrial vision setup.

Results & Evaluation

- YOLOv11n achieved the strongest overall detection performance across the dataset sizes tested.
- YOLOv8n performed well and provided reliable real-time detection during edge deployment.
- YOLOv5n showed moderate performance while PP-PicoDet had the lowest overall accuracy.
- Overall, detection performance declined as the training dataset was reduced, with smaller models showing greater sensitivity to limited data.
- The industrial vision comparison was limited by the use of a monochrome Cognex camera, which restricted colour-based testing on the fruit dataset.

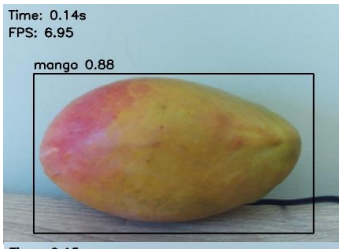

YOLOv11n Performance

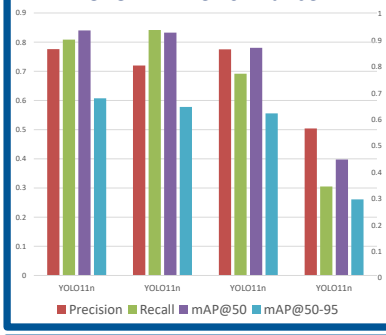

YOLOv8n Performance

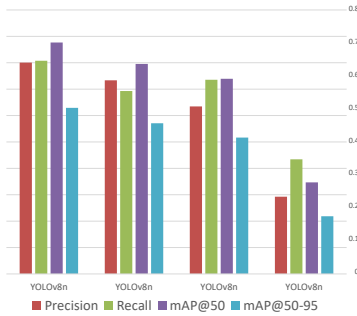

YOLOv5n Performance

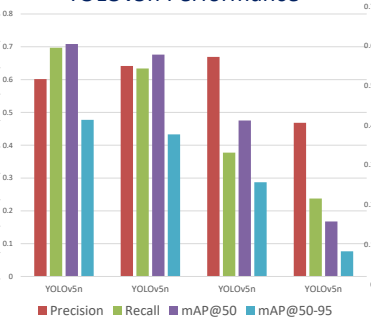

PP-PicoDet Performance

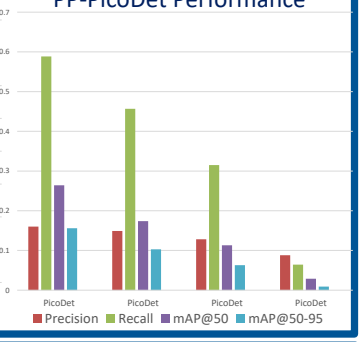

Future Work & Conclusion

- The project demonstrated that lightweight object detection models can support real-time fruit classification on edge devices.
- Performance depended not only on model architecture but also on the amount of training data available.
- Future work will extend testing to more challenging images conditions including changes in lighting, background and fruit placement.
- Future industrial evaluation will use a colour industrial camera to allow a more direct comparison with the Edge AI models.

Acknowledgments

I would like to thank Ted Scully and Collin Manning for their supervision and Dell Technologies for sponsoring this project.

IT & Security Learning Web-Platform SecuLab

Mykhailo Budonnyi,
BSc Honours in IT Management and Cybersecurity
Department of Computer Science,
MTU Cork, May 2026

Introduction

Cybersecurity is critical part of digital environment where human error remains the leading cause of security breaches. Increasing attack surfaces (e.g., cloud, remote work) expose users to daily threats. It is necessary to educate users and use effective learning methods which be **interactive and scenario-based experience** rather than passive theory.

Problem Statement

How can traditional cybersecurity training be improved to become more engaging and effective in developing real-world security skills?

Proposed Platform

SECULAB is an interactive, gamified cybersecurity training platform that uses real-world scenarios to improve user engagement and practical decision-making skills.

Scenario Engine

Core Modules

- Scenario-Based Learning
- Gamification
- User Progress Tracking
- Real-time feedback

Prototype

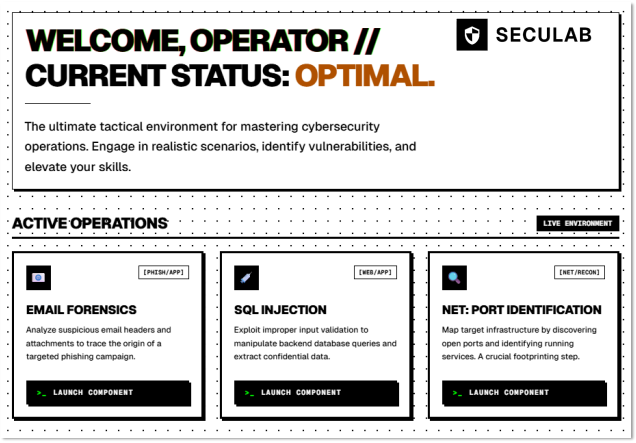

Platform Architecture

Interface → Application → Gamification → Database

Results & Conclusion

Cybersecurity Knowledge Level of Test Users

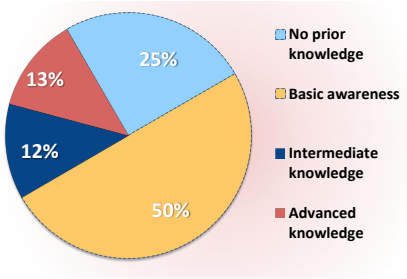

Most Effective Learning Components Selected by Test Users

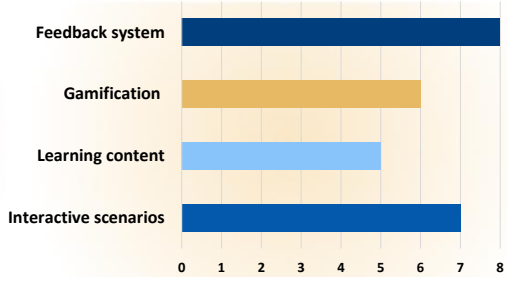

Improvement in User Confidence after Using SecuLab

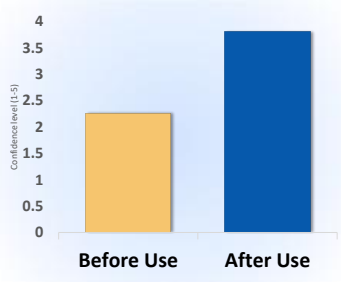

- SecuLab demonstrates that gamified, scenario-based learning **enhances user engagement** and **supports cybersecurity awareness** for a broad, non-technical audience.
- User testing indicates improved confidence, high engagement, and positive feedback on scenario-based learning.

Acknowledgments

The author would like to thank Larkin Cunningham and Kashif Ahmad for their supervision, guidance and support during this project

Can lecture slides become reliable quiz questions without hallucinating?

AutoQuiz

Automatic MCQ Generation for Computer Science Lecture Slides

Karim Cali · BSc Honours in Software Development · Department of Computer Science, MTU Cork · May 2026

① The Problem

Students revise mainly from lecture slides — but that revision is **passive and slow.**

Passive reading
Re-reading slides
no active recall
forgotten quickly

→

Active retrieval
Answer questions
test yourself
retain far more

Research shows retrieval practice outperforms passive re-reading for long-term memory.

The gap:
most quiz generators work on long clean text — not on bullet-heavy slides, tables, and diagrams.

② Why Not Just Use LLMs?

Six leading LLMs were tested on the same CS slides (Nov 2025). All wrote fluent MCQs - but failed at the one thing that matters:

Questions weren't grounded in the actual slides

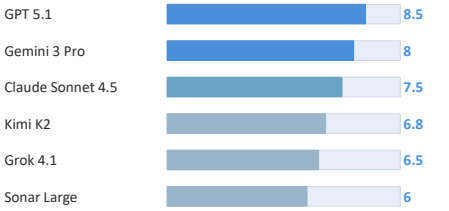

Common gaps across ALL models:

✗ Questions not tied to specific slide text
✗ Topic coverage didn't reflect the actual deck
✗ No difficulty control — mostly basic recall
✗ No metadata: no slide reference, no source

③ Challenges

Parsing Messy Slides
Some PDF and PPTX decks had broken text, admin content, and weak fragments.
Fix: add content filtering, grouping, and usable slide selection.

Distractor Quality
Broad wrong answers reduced question quality
Fix: build distractors from a curated cs lexicon and slide-local terms.

Duplicate Control
Similar slide content could produce repeated stems.
Fix: Score drafts for overlap, clarity, and answerability before delivery.

OCR cleanup
Screenshot-heavy slides produced both useful OCR and noisy OCR.
Fix: keep readable OCR blocks and reject noisy text

④ How AutoQuiz Works

Upload slides → get a quiz. Every question traces back to a specific slide.

1. **Upload** — PDF / PPTX lecture decks
2. **Parse** — Extract headings, nested bullets (levels 0–2), tables (headers+cells), figures + captions
3. **Confidence** — H = strong structure · M = usable · L = skipped. Only H+M reach generation
4. **Templates** — MCQ, cloze, table-lookup — rule-based, no AI
5. **Review** — Auto-check: answerability, clarity, duplicates
6. **Quiz** — Balanced 20-item quiz + slide-grounded feedback

⑤ CS Lexicon & Distractors

Good MCQs need wrong answers that are **plausible but clearly incorrect.** A curated CS lexicon provides:

Siblings
Closely related concepts
in the same domain
e.g. TCP → UDP, HTTP, FTP

Near-miss
Plausible confusions
from the same lecture
e.g. TCP → socket, packet, IP

Domains covered: Software Engineering · Networks · Databases · OS · AI

Quiz Policy

Normal	Short	Min
20 items	10 items	8 items

⑥ Worked Example: Slide → Question

Here's exactly how AutoQuiz turns a slide bullet into an MCQ:

SLIDE CONTENT

"TCP is a connection-oriented protocol."

↓

KEY TERM: **TCP**

DISTRACTORS (from lexicon): UDP · HTTP · FTP

↓

GENERATED MCQ

Which protocol is connection-oriented?

- A TCP ✓
- B UDP
- C HTTP
- D FTP

Source: Slide 12, Bullet 1 · Template: definition · Confidence: High

⑦ What the System Supports

PDF & PPTX ingestion
Validates files; detects encrypted or image-only decks

Structured parsing
Headings, nested bullets, tables, figures — with H/M/L confidence tags

6 question templates
nested_list · section_list · flat_list · definition · table_lookup · table_max

Smart distractor generation
CS lexicon siblings + slide-local near-miss terms

Automated review rubric
Answerability · clarity · distractor quality · near-duplicate detection

Student-first quiz flow
Upload → Start Quiz — no editing required; balanced 20-item default

Staff review mode
Accept · edit · reject with reason codes; optional Canvas/QTI export

⑧ The 6 Question Templates

Rule-based only — no AI. Every question traces back to a specific slide.

nested_list
Multi-select
Used for slides with main points and sub-points.
e.g. Which of the following are strengths of Scrum?

section_list
Multi-select
Used for slides split into labelled sections such as features, uses, or steps.
e.g. Which describe how to interact with Git?

flat_list
MCQ
Used for short bullet lists with one clear correct item.
e.g. Which is a limitation of Waterfall?

definition
MCQ
Used for direct fact or term-definition statements.
e.g. Which option best defines TCP?

table_lookup
MCQ
Used when the answer comes from one table value.
e.g. Which protocol uses port 443

table_max
MCQ
Used when the question asks for the highest or lowest value in a table.
e.g. Which tool has the highest accuracy?

⑨ Conclusions & Future Work

✓ Deterministic templates produce slide-grounded, traceable MCQs

✓ LLMs score well — but can't guarantee slide-specific questions

✓ CS lexicon enables high-quality, plausible wrong answers

✓ Automated review filters weak items before students see them

Future work:

OCR for image-heavy slides · Canvas/QTI export · Adaptive question ordering · Broader domain lexicons

Acknowledgments

Thanks to Alex Vakaloudis for supervision and guidance throughout the research and prototype stages of this project, and to the MTU Department of Computer Science.

References

[1] J. Dunlosky et al., "Improving students' learning with effective techniques," Psychol. Sci. Public Interest, 2013.

[2] G. Laban et al., "Quiz design task: teacher-in-the-loop question generation," arXiv, 2022.

[3] S. Maheen et al., "Automatic CS MCQ generation from informative sentences," PeerJ CS, 2022.

[4] S. Mulla & P. Gharpure, "Automatic question generation: a review," Prog. Artif. Intell., 2023.

[5] H. L. Roediger & J. D. Karpicke, "Test-enhanced learning," Psychol. Sci., 2006.

Takeaway: The strongest path isn't "let an LLM invent quiz questions." It's turning **slide structure into explainable, traceable questions** grounded in how the module is actually taught.

A Degree Planning and Progress Tracking Application for Cyberskills Ireland students

MTU
Ollscoil Teicneolaíochta na Mumhan
Munster Technological University

Daniel Callanan, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Many Cyberskills Ireland students build their qualifications over several semesters or years. This flexible study pattern gives students more choice, but it can make progress hard to track. Module records, credit totals, and award rules often sit across emails, handbooks, and personal notes. This creates confusion and makes long term planning harder. This project proposes a web application that brings these records into one clear system. It shows completed modules, earned credits, and possible award paths in a simple format. It uses dashboards and a chat assistant to help students check progress, understand requirements, and plan future study with confidence.

Architecture

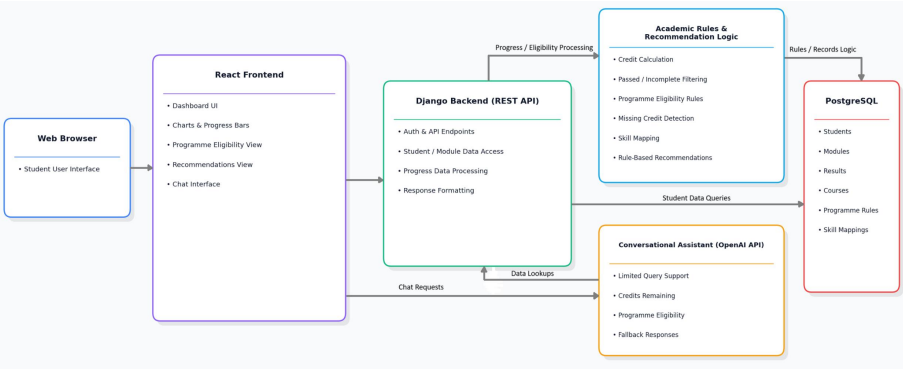

Project Goal

The goal of this project is to design and build a web application for degree planning and progress tracking. The system will give Cyberskills students one place to review their academic record and plan next steps. Students will see completed modules, total credits, remaining requirements, and eligible programmes. A visual dashboard and chat assistant will present this data clearly and support quick questions about progress and study options.

Application

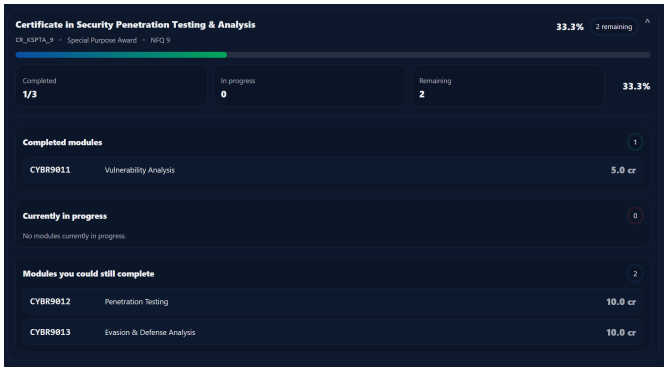

Programme Eligibility Tracking

Student can view:

- Completed modules
- Modules in progress
- Module options to complete programme

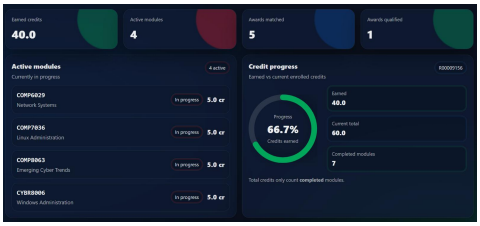

Main Dashboard

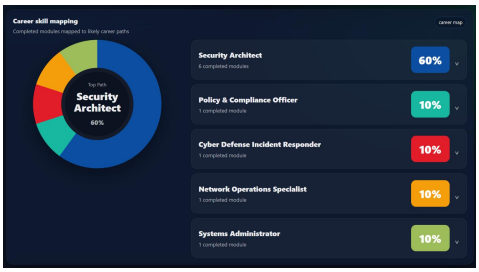

Modules Completed mapped to career paths

Technologies used

Acknowledgments: I would like to thank Jacqueline Kehoe for her support and guidance throughout this project.

Readmit Predictor:

Software System for Correlating Clinical Language with 30-Day Patient Readmissions

Seán Collins, BSc Honours in Software Development
Department of Computer Science, MTU Cork, May 2026

Research Question & Project Goal

Can a correlation be made between the language found in medical staff plain text notes and the frequency of re-admittance of a patient within a time-span of thirty days post discharge;

This project aims to see if there is predictive potential in the discharge summaries of patients on whether they would return in 30 days and provide an application to easily provide prediction feedback from patient records.

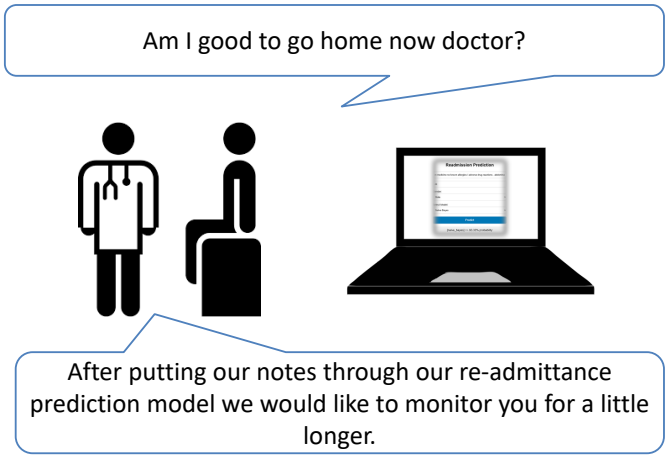

Data Acquisition

Models were trained on the MIMIC-IV dataset, which is an anonymised dataset of patient records from the **Beth Israel Deaconess Medical Centre (BIDMC)** in Boston, Massachusetts, acquired from the PhysioNet site upon completion of an ethical use of data course.

Modelling Approach

Selected algorithms to test include:

Naïve Bayes **P(c|d) = (P(d|c) * P(c)) / P(d)**

Logical Regression **σ(z) = 1 / (1 + exp(-z))**

SVM (Support Vector Machine)

Tech. Used

Implementation

Frontend Application

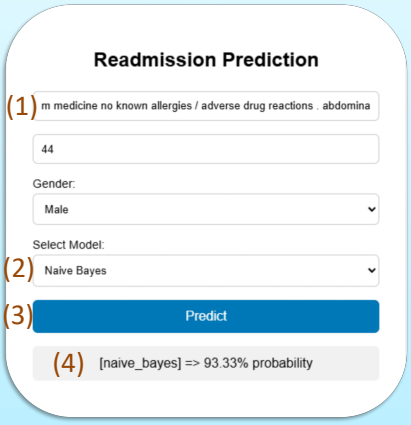

(1). Input fields to enter discharge summary, patient age and patient sex.
(2). Drop down select an algorithm to apply data to.
(3). Button when all data is entered to run through the model
(4). Output showing percentage likelihood of patient readmission.

Evaluation

NLP Algorithm Comparisons

Confusion Matrices

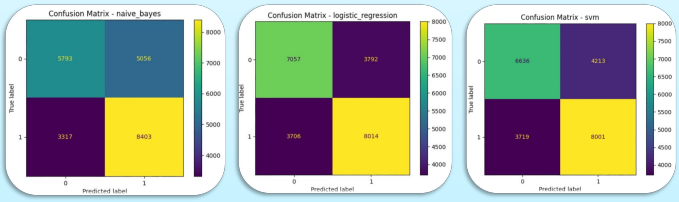

Comparisons

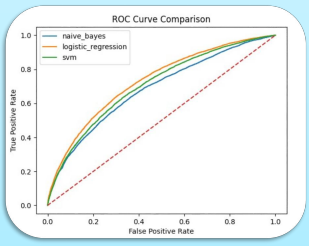

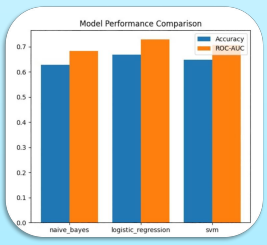

	Accuracy	ROC-AUC
Naïve Bayes:	*0.629*	*0.682*
Logistical Regression:	*0.668*	*0.729*
SVM:	*0.649*	*0.706*

Conclusions

After looking at the predictive ability of discharge notes using these models shows there's a potential in predicting remittance however at this moment also applying structured data proved to provided a stronger predictive potential.

References

1. Johnson, A. E. W., Pollard, T. J., Shen, L., et al., "MIMIC-IV (version 2.2): A freely accessible electronic health record dataset", PhysioNet, 2023.Etc.

Acknowledgments

This author would like to acknowledge Colin E. Manning for providing advice and assistance as project supervisor.

Hollow Signal

A Cosmic Horror Video Game

Ciara Dalton, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

The video game industry has grown significantly in cultural impact, with horror games proving especially effective at creating immersive emotional experiences. While simulated desktop interfaces have been used as narrative devices in games like *Her Story* and *Hypnospace Outlaw*, their application within cosmic horror remains unexplored.
Hollow Signal addresses this gap by developing a fully diegetic cosmic horror game where all narrative and horror emerge through interaction with a simulated early 2000s computer interface.

Game Overview

Hollow Signal is a cosmic horror game played entirely through a simulated early 2000s desktop. The player takes the role of Sam Carpenter, who gains access to the personal computer of their missing friend Ellie Flanagan.
By exploring files, diary entries, emails, and online forums, the player gradually uncovers the truth and the presence of something not of this world.

Key Features

- Simulated OS with hierarchical file system and interactive applications
- State management system tracking player actions and narrative progression
- Password/access mechanics as investigative gameplay puzzles
- Horror event system triggering visual glitches and audio distortion in response to game state

Architecture

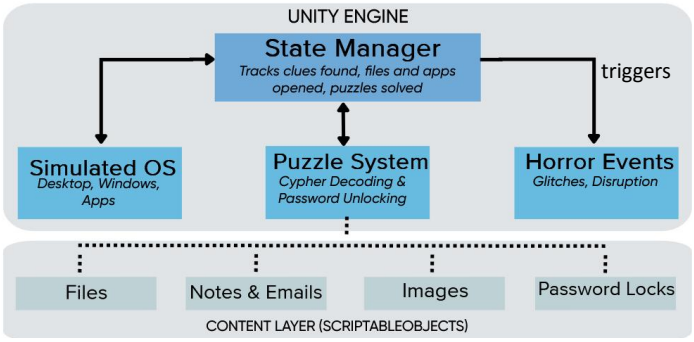

A Deeper Look

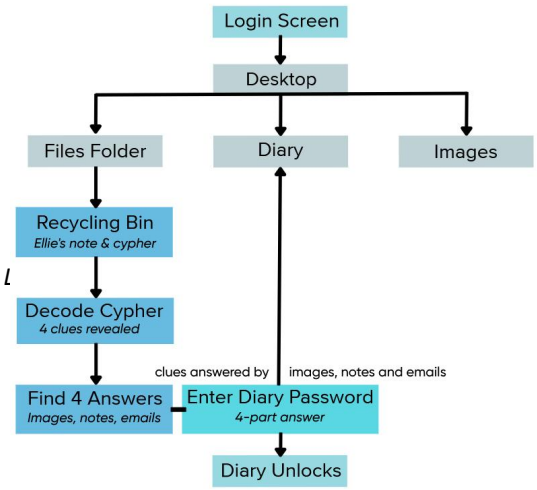

A flow diagram showing the path of the 'vertical slice' demo developed during the implementation phase

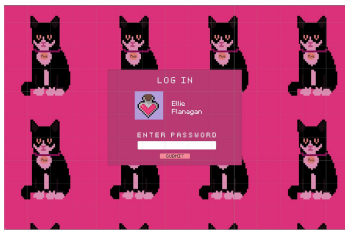

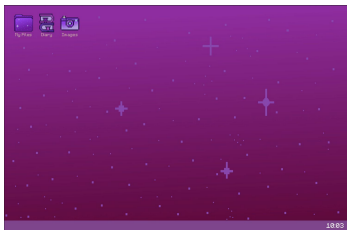

Login and desktop scene

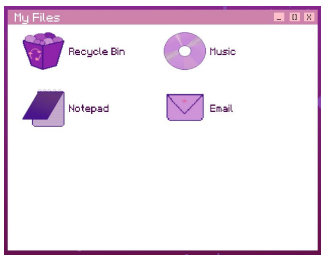

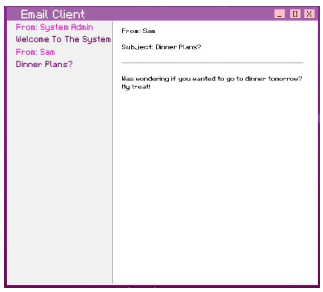

Files folder and email application

Conclusions

Hollow Signal demonstrates that a simulated desktop interface can function as an effective game world, addressing a gap at the intersection of cosmic horror and diegetic interface design. The project required designing a modular Unity architecture supporting investigative gameplay, state management across multiple systems, and responsive horror event triggering. Future work includes completing the full narrative, expanding the puzzle system, and developing the full game.

Acknowledgments

The author would like to acknowledge her supervisor Larkin Cunningham for his invaluable support and advice.

Website for Hotel Bookings

Nikola Denisova, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Many online hotel booking platforms suffer from usability issues such as hidden fees, difficult interfaces and repetitive search processes. This can reduce user trust and negatively impact the booking service.

Problem Statement

Current hotel booking systems:

- ❑ Misleading prices – excluding taxes and fees.
- ❑ Users must repeatedly apply the same filters.

Challenge:

Develop a website which aims to address these problems through a more clear and user-focused design.

Objective

- ❑ Develop a Java-based web application for hotel bookings.
- ❑ Provide transparent total pricing upfront.
- ❑ Allow users to save and reuse search filters.
- ❑ Implement personalised hotel recommendations.
- ❑ Improve user experience and usability.

Methodology

- ❑ Research existing hotel booking platforms.
- ❑ Identify usability and transparency issues.
- ❑ Design system architecture and database.
- ❑ Test functionality and user experience.

Key Features

- ❑ Transparent Pricing.
- ❑ Saved Filters.
- ❑ Smart Recommendations.
- ❑ Booking Systems.

Results

- ❑ Improved user trust through price transparency.
- ❑ Provides a more personalised user experience.
- ❑ Demonstrates a functional web application.

Software Architecture

The system follows a layered web app architecture:

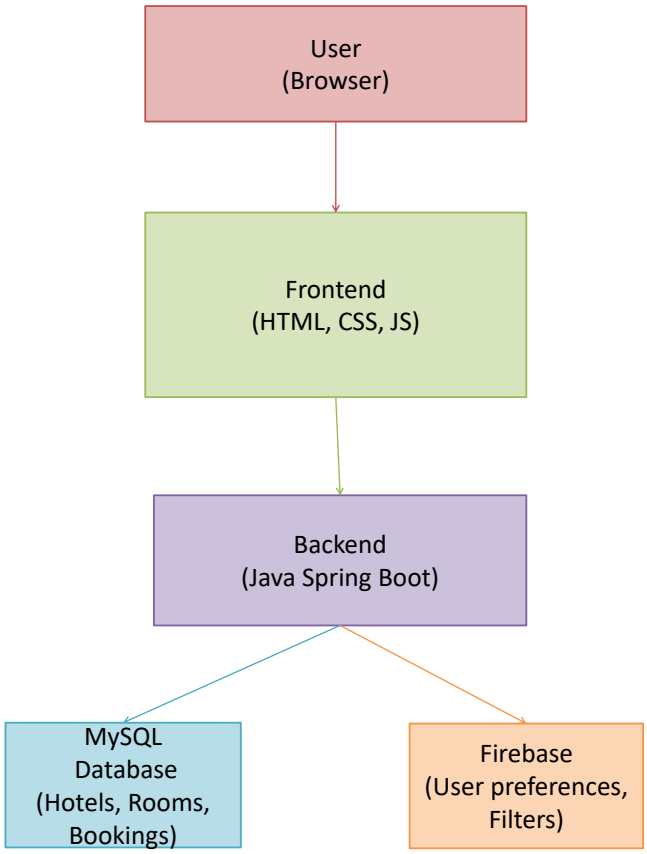

Conclusions

This project demonstrates how targeted improvements in transparency and usability can significantly enhance the hotel booking experience. The systems shows that effective personalisation can be achieved without complex AI, using simple data-driven techniques. The final application provides a user-focused alternative to existing booking platforms.

Acknowledgments

I would like to thank both Mohammed Hasanuzzaman and Colin Manning for supervising this project and my classmates for the constant support.

Leveraging Supervised Learning to Detect and Classify Misinformation in Political Media

Morgan Flynn, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

The proliferation of AI-generated political articles has plagued various online media platforms, raising concerns for their reliability and credibility. This observation led to an innovative thought: leveraging supervised learning algorithms as an approach to addressing this issue. Machine Learning is inherently associated with solving numerous detection challenges, particularly in textual contexts.

Multimodal Solution

The proposed framework uses a binary classification approach to raise widespread awareness of misinformative content in media. By leveraging supervised learning, we meticulously analyse text data to identify patterns indicative of informative or misinformative content. This is further extended by incorporating a pretrained Optical Character Recognition(OCR) model, EasyOCR, to extract and analyse text embedded within images such as screenshots and infographics, creating a multimodal solution that promotes transparency and accountability in political digital media.

AI Pipeline

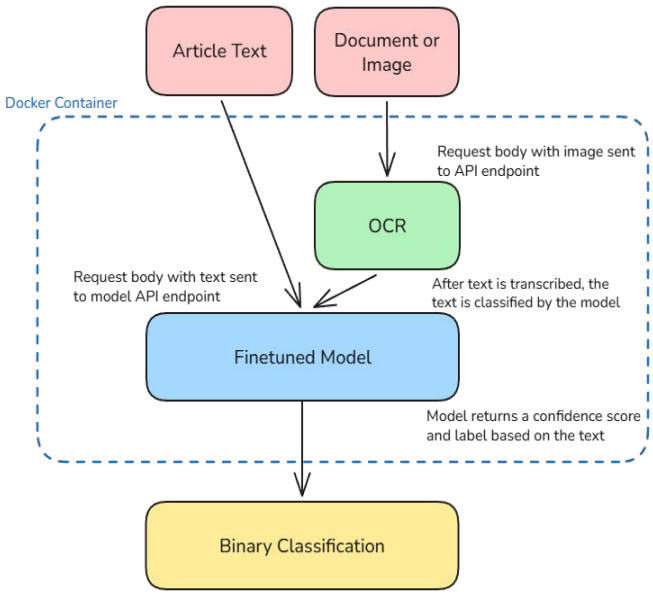

The pipeline accepts two input modalities: raw text and image-based content. Images are first processed by EasyOCR to extract readable text, before both pathways converge on the finetuned DistilBERT model, which returns a binary classification and confidence score.

Methodology

config	learning_rate	num_train_epochs	per_device_train_batch_size	weight_decay
kaggle_baseline	5.00E-05	3	16	0.01
welfake_lr_low	2.00E-05	3	16	0.01
welfake_lr_high	0.0001	3	16	0.01
welfake_epochs_long	5.00E-05	5	16	0.01
welfake_bs_large	5.00E-05	3	32	0.01
welfake_wd_none	5.00E-05	3	16	0
welfake_wd_high	5.00E-05	3	16	0.1
welfake_bs_small	5.00E-05	3	8	0.01
welfake_epochs_shor	5.00E-05	1	16	0.01

config	accuracy	f1	precision
kaggle_baseline	0.988995441	0.9889934753	0.9889986127
welfake_lr_low	0.9910391448	0.9910374078	0.9910440251
welfake_lr_high	0.988366609	0.9883638194	0.9883736524
welfake_epochs_long	0.989152649	0.9891515283	0.9891528244
welfake_bs_large	0.9897814809	0.9897791876	0.9897880039
welfake_wd_none	0.9915107687	0.9915102702	0.9915105625
welfake_wd_high	0.9905675208	0.9905675208	0.9905675208
welfake_bs_small	0.9902531049	0.9902522432	0.990253159
welfake_epochs_short	0.9893098569	0.9893089119	0.9893097426

DistilBERT

Model Confusion Matrix

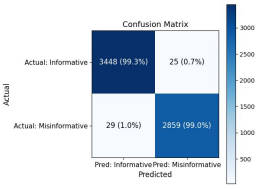

OCR

EasyOCR

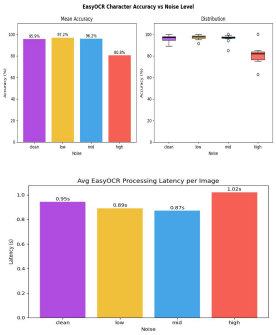

UI

Final Design

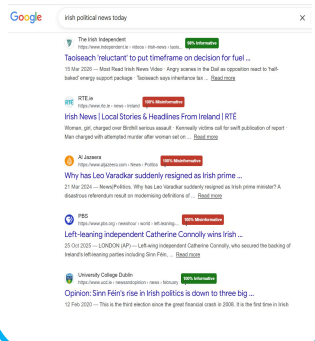

Technologies Used

Conclusions

- **The system demonstrates strong training performance (~0.98 accuracy), but is unrealistic for a deployed model considering the complexity of detecting linguistic nuance in live data (~0.50 accuracy).**
- **Multimodal analysis of both text and imagery broadens detection coverage beyond text-only approaches.**
- **Classifiers trained on historically labelled news data cannot reliably generalise to emerging misinformation strategies, outlining the need for continuous retraining as the political media landscape evolves.**
- **Satirical content may occasionally be flagged as misinformative, underscoring the value of human judgement as a final check.**

Smart Touchless Control and Home Automation IoT System Using Embedded Sensors and Cloud Integration

Luis Giliberti, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

INTRODUCTION

Smart home systems allow users to control devices through mobile apps and voice control.
This project introduces a touchless gesture based control system using embedded sensors, an ESP32-C6, and cloud integration. It enables intuitive physical controls with real-time response, while maintaining cloud connectivity and mobile app access.

PROBLEM STATEMENT

Current smart home systems:
Require physical interaction (unhygienic)
Advanced gesture systems are expensive and complex
Limited control methods (app or voice only)
Challenge:
Design a system that enables fast, reliable, and touchless control with both edge and cloud capabilities.

OBJECTIVE

- Detect gestures using ESP32-C6 + ToF sensors
- Provide real-time visual feedback
- Sync data using Firebase Firestore
- Enable control via Android mobile app
- Support Node.js backend integration
- Ensure low latency

SOFTWARE & HARDWARE ARCHITECTURE

FLOWCHART

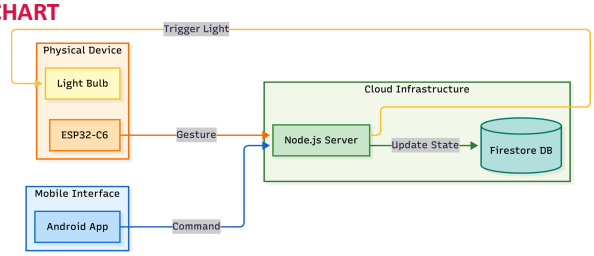

SEQUENCE DIAGRAM

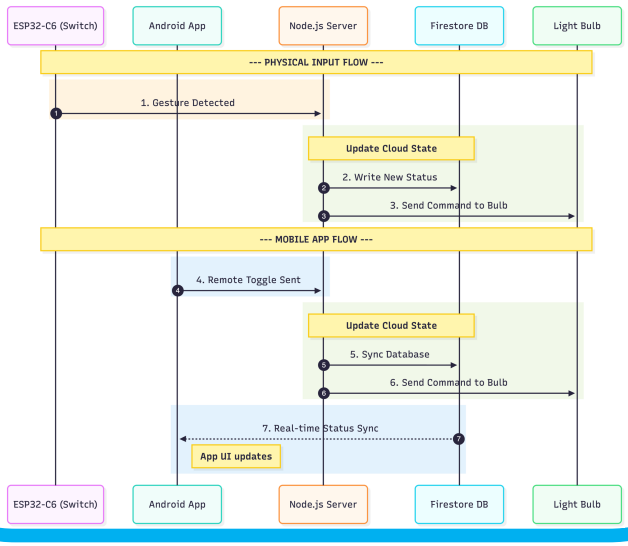

METHODOLOGY

- The system uses Time-of-Flight sensors to detect gestures.
- The ESP32 processes the data locally and provides instant feedback via TFT display.
- Events are sent to a Node.js server, stored in Firestore, and displayed in the Android app.

RESULTS

- Real-time gesture detection achieved
- System latency avg ≈ 400–700 ms
- Reliable cloud synchronization
- Functional Android control dashboard
- Successful smart light automation

HARDWARE

XIAO ESP32C6

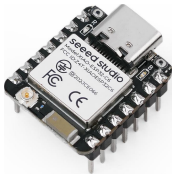

VL53L0XV2 ToF Sensor

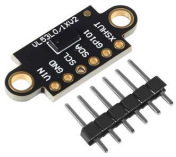

TFT DISPLAY

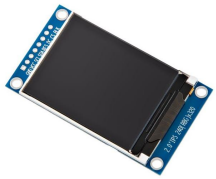

SMART LIGHT BULB

GRAPHS

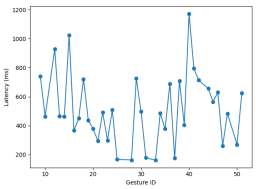

Time taken from when a gesture is detected on the ESP32 to when the server receives the request.

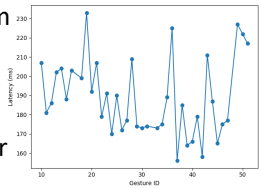

Time taken from when the server processes the request to light bulb updates.

Conclusions

The project successfully demonstrates a touchless smart home control system combining: Embedded systems Cloud computing Mobile applications The system achieves low latency, real-time feedback, and remote control, but future improvements are needed for accuracy and gesture recognition.

Acknowledgments Thanks to my supervisor and classmates who supported the development of this project.

The Value of Traditional vs Contemporary Game Development

Matthew Greaney, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Topic Introduction & Research Question

Game development has changed dramatically from the early arcade and console titles to today's games pushed forward by powerful engine workflows. This project explores those changes by developing a single game using both **traditional** and **contemporary** development approaches. The project focuses on how these different methodologies affect the **development experience** and the **final game**.

Research Question
Do traditional gaming development methods offer comparable educational and practical value to contemporary approaches?

Methodology

Tools and Technology

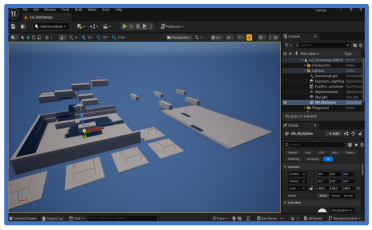

Unreal Engine Editor

A widely adopted and used modern game engine and editing suite

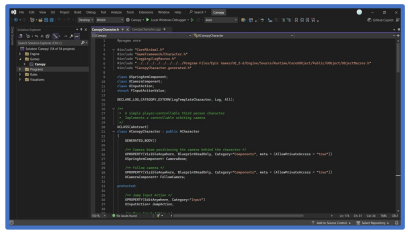

Microsoft Visual Studio

A powerful IDE from Microsoft that lets developers write, debug, test, and deploy software

Project Workflow

- Research and Planning
- Prototyping
- Implementation
- Polish
- Reflection

Traditional & Contemporary Work in Practice

Traditional Elements:
- Character Movement System
- Input Handling
- Character Stats

Contemporary Elements:
- Level Design
- On Screen HUD
- Sound and Music
- VFX

Results and Discussion

Contemporary tools made development noticeably faster and easier, especially for level building, debugging and polish, but many systems were heavily abstracted. In contrast, developing aspects in a more traditional, low level way took longer but gave much finer control and significantly deepened my understanding of core game programming concepts.

Overall, the work suggests that contemporary methods are more efficient, while traditional methods remain highly valuable for building strong technical insight.

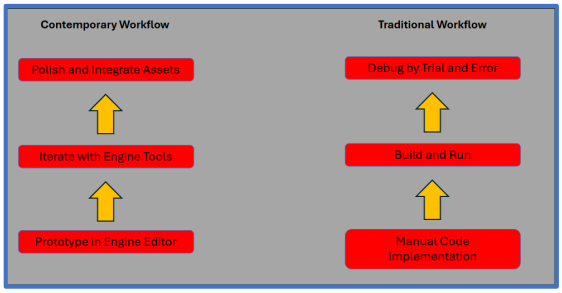

Important References

1. Kin W., "Game Engines: Revolutionizing Digital Creativity Across Industries", International Journal of Advancements in Technology, 2024

Acknowledgments

I would like to acknowledge Dr.Larkin Cunningham as the discussions that were had about the industry informed this project greatly. I would also like to acknowledge my peers Ryan and Jason for testing.

Design and Evaluation of a Gamified Cybersecurity Awareness Platform for Improving User Engagement

Muhammad Haseeb, BSc Honours in IT Management
Department of Computer Science,
MTU Cork, May 2026

Introduction

- Phishing is one of the leading causes of data breaches.
- Humans remain the weakest link despite advances in security technologies.
- Traditional cybersecurity training is:
  - ❑ Passive
  - ❑ Tedious
  - ❑ Ineffective

➢ **This project investigates whether gamification can improve user engagement and learning.**

Objectives

➢ Develop a gamified cybersecurity training platform
➢ Teach users to identify phishing attacks and common threats
➢ Improve engagement using game mechanics
➢ Evaluate:
  - ❑ Learning effectiveness
  - ❑ User engagement
  - ❑ Usability

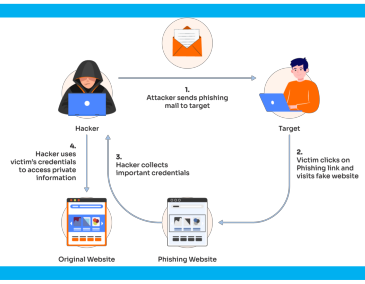

System Design

Architecture Overview:

- Frontend: Game interface (interactive learning)
- Backend: Logic, scoring and progression
- Database: Stores user progress and performance

Game Flow:
Play Scenario → Make Decision → Get Feedback → Earn XP → Progress

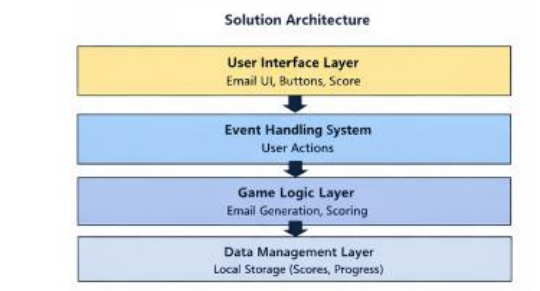

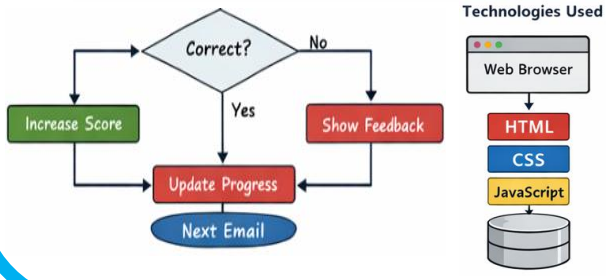

Key Features

🎯 Phishing Detection Scenarios

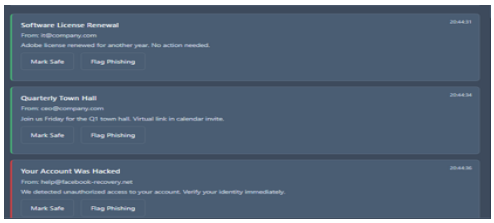

📊 Score and Progress System

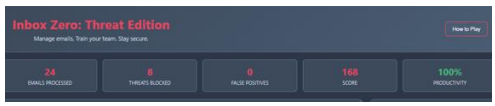

🕹️ Game Dashboard (Score, levels, feedback, automation)

🔁 Replayability and Feedback

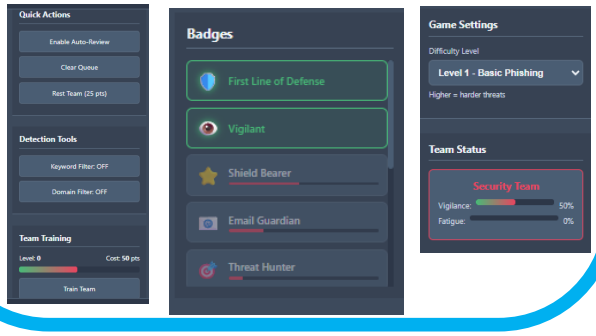

Results & Findings

- Gamification can significantly improve user engagement
- Interactive learning improves knowledge retention
- Immediate feedback reinforces learning

➢ **Gamified learning is more effective than traditional methods for cybersecurity awareness.**

References

1. J. Koivisto and J. Hamari, "The rise of motivational information systems: A review of gamification research," International Journal of Information Management, vol. 45, pp. 191–210, Apr. 2019.
2. R. Kalakuntla, A. B. Vanamala, and R. R. Kolipyaka, "Cyber Security," HOLISTICA – Journal of Business and Public Administration, vol. 10, no. 2, pp. 115–128, 2019.

Acknowledgments

The author would like to acknowledge the guidance provided by his supervisor Larkin Cunningham.

Detecting Fraudulent Transactions using Ensemble Learning with Explainability

Kevin Irabor, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

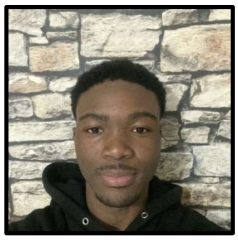

Introduction

Ethereum processes billions in transactions daily, yet its pseudonymous architecture enables large-scale financial crime. In **2024**, illicit cryptocurrency activity exceeded **$40.1 billion** globally.
Rule-based fraud filters are easily evaded as attackers adapt behaviour

- Transactions are irreversible — missed fraud cannot be recovered
- Objective: detect fraudulent Ethereum accounts using ensemble ML and explain model decisions using SHAP

Research Question

RQ1: Do ensemble models outperform single classifiers for Ethereum fraud detection?
RQ2: How does SMOTE affect fraud detection recall?
RQ3: Which features most influence fraud classification (SHAP)?
RQ4: Are SHAP rankings consistent across model families?

Methodology

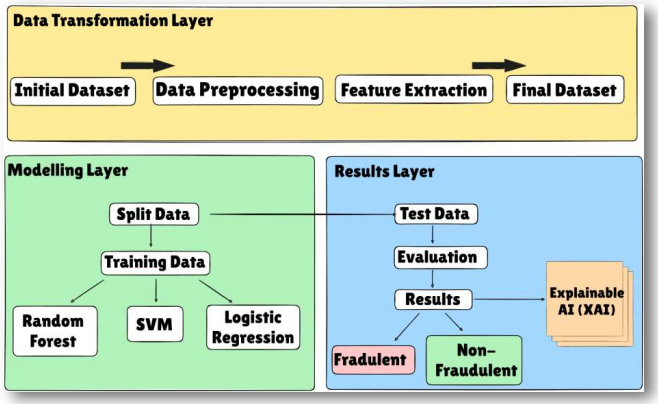

Results

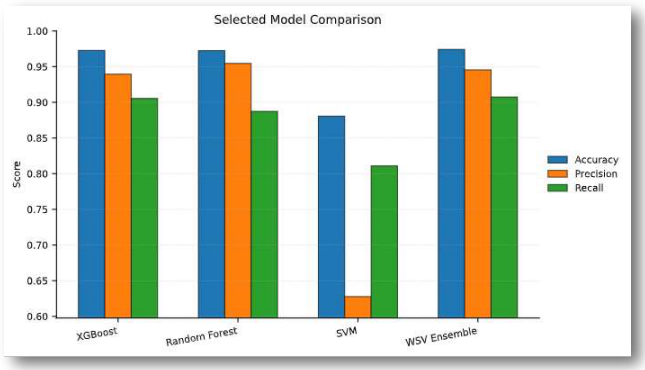

SHAP Explainability

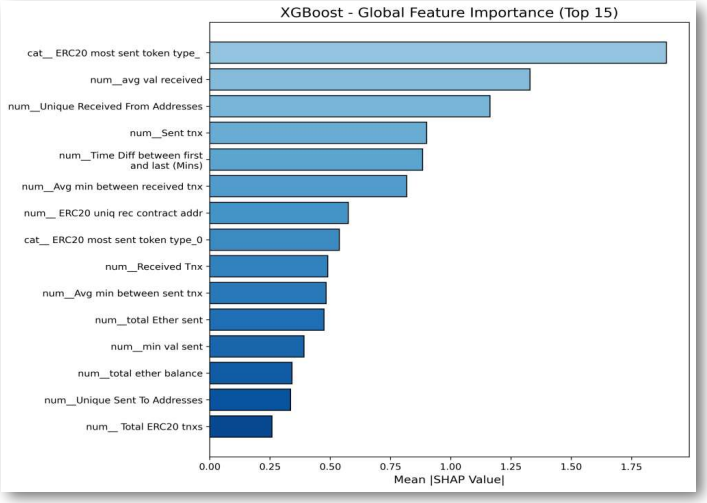

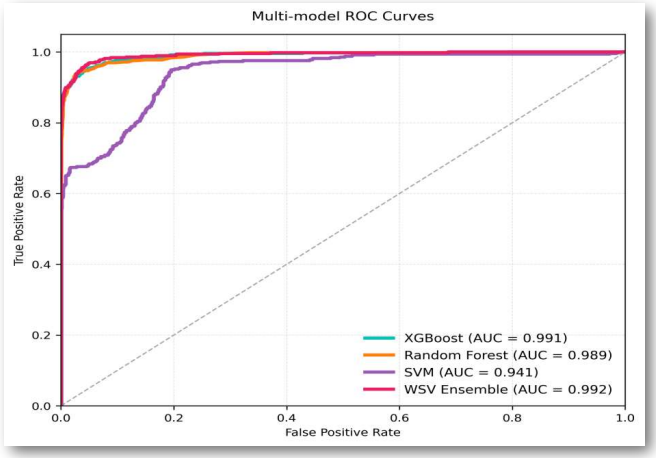

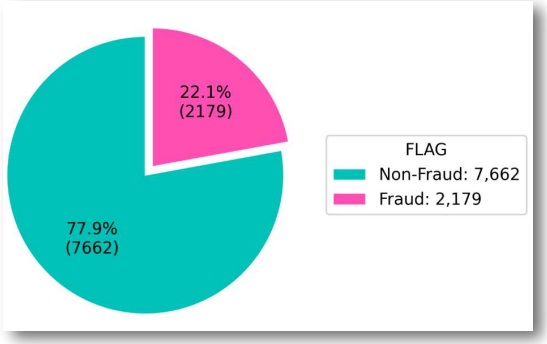

Conclusion

- WSV & XGBoost (91.8%, 90.5% recall, AUC ≥ 0.991) outperform SVM baseline (81.1%)
- ERC20 token identity dominates SHAP mandatory over Gain
- Fraudster evasion via legitimate mimicry - future GNN approaches

Future Work

- GNNs for fraud ring detection
- Real-time Google BigQuery integration
- Cross-fold SHAP stability analysis

GitHub

References

1. Chainalysis, "Crypto Crime Report," 2025.
2. Vagifa, "Ethereum Fraud Detection Dataset," Kaggle, 2021.
3. S. Farrugia et al., "Detection of illicit accounts," Expert Systems, 2020.

Acknowledgments

Special thanks to my supervisor Jing Hua Ye, Ruairí Ó Reilly, my colleagues and my family for support and guidance throughout the project.

Web Application To Build Lightweight Responsive Websites

Jack Jackson-Lane,
BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

In today's age, having an online presence is more important than ever. Whether you are a business looking to grow, a writer wanting to share your ideas or a student wanting to showcase your personal work, this web application solves this problem by providing a platform to build lightweight and responsive websites for free without needing any coding experience.

Objective

The objective of this project is to provide a platform that helps users create lightweight, responsive websites.

Using this application, users can design their own sites and download them to their local machine. A variety of templates are available to help users get started.

This project also provides hosting tips to help users host their website.

Application

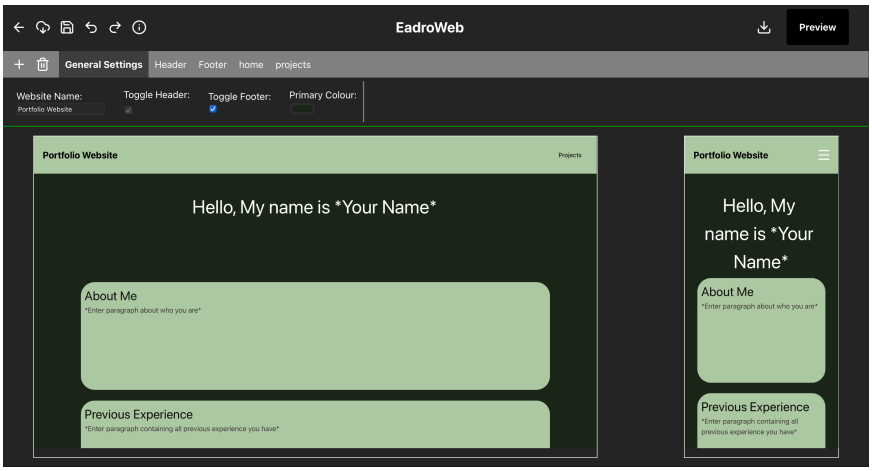

Customisation Page

Architecture

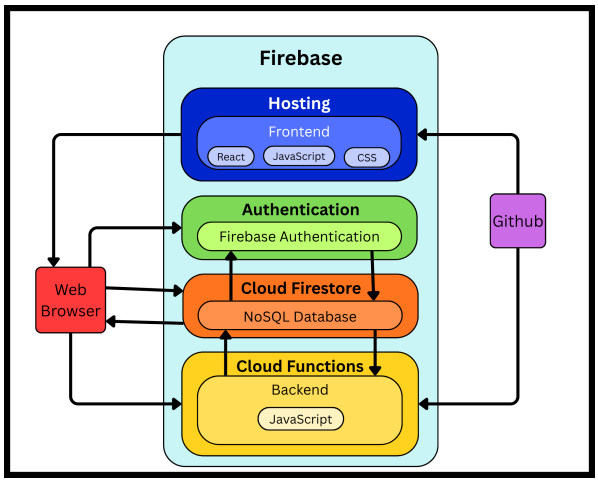

Downloaded Website

Technologies

Firebase App Hosting
Firebase Cloud Functions for Firebase
Firebase Authentication
Firebase Cloud Firestore

Acknowledgments

I would like to thank my supervisor, Irene Foley for her continuous support and guidance with this project.

Continuous Authentication via Gesture, Pressure and Timing Patterns on Android

Joshua Jones, BSc Honours in Computer Systems
Department of Computer Science,
MTU Cork, May 2026

Introduction

Modern smartphones rely on a one-time authentication such as PINs or fingerprint scans. Once a device is unlocked, anyone holding it has full access with no further checks. Continuous authentication solves this by verifying the user throughout their normal use. Behavioural biometrics capture the unique way a person swipes and presses the screen and are difficult to copy even when the shape of a gesture is known to an attacker. This project builds a gesture-based continuous authentication system for Android using only public APIs, no root access and fully on-device storage to protect user privacy.

System Architecture and Feature Set

The system follows a five-stage modular pipeline: **Gesture Capture** ⇨ **Feature Extraction** ⇨ **Behavioural Profile** ⇨ **Classifier** ⇨ **Auth Decision**

During registration, the user draws a gesture repeatedly to build a statistical profile. During authentication, each new gesture is compared using Euclidean distance and per-feature z-score checks. A dynamic threshold tightens as the profile grows from 4.5 at under 100 samples down to 2.5 at over 400, so accuracy improves over time. The system extracts 15 behavioural features from raw touch data: gesture duration, mean and peak velocity, velocity variance, stroke length, directness, curvature, inter-event timing, timing variance, device pitch and roll and four pressure features. Pressure is kept but given a reduced weight of 0.1 as the test device uses a capacitive screen that estimates rather than directly measures force.

Results and Evaluation

False Rejection Rate (FRR)
FRR = rejected / total genuine

The proportion of genuine user attempts that the system incorrectly rejects. A high FRR means the real user is blocked too often.

False Acceptance Rate (FAR)
FAR = accepted / total impostor

The proportion of impostor attempts the system incorrectly accepts. A high FAR means the system is too lenient with attackers.

72.7 µs
Mean classification time (target was 50 to 100ms)

80%
Accuracy under ideal conditions

75% Average
confidence score

Condition	Accepted	FRR	Avg Conf
Right hand, seated	8 / 10	20%	79%
Left hand, seated	5 / 10	50%	69%
Right hand, standing	5 / 10	50%	73%
Left hand, standing	7 / 10	30%	80%
Overall	25 / 40	37.5%	75%

Profile size: 100 to 200 samples. Higher FRR in cross-condition tests is because of the profile being built from right-hand seated gestures only. Accuracy should improve as the profile increases toward 400 universal samples (all sorts of conditions)

Bonus Security Modules

Extending the System Beyond In-App Gestures

Android sandboxing prevents reading touch events outside the application without root access, so gesture monitoring is limited to within the app. Two background modules compensate for this.

Jerk and Snatch Detection is a background service that checks the accelerometer continuously. When it detects a sharp physical acceleration above 700 metres per second cubed (m/s³), such as a phone being grabbed, it triggers an immediate device lock through the Android DevicePolicyManager. It runs regardless of which app is in the foreground.

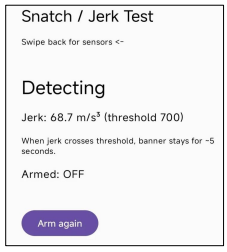

Location Trust Scoring is an hourly WorkManager background job. It tracks how often the device visits different areas and marks locations visited three or more times per week as trusted. Unfamiliar locations signal a higher-risk context for authentication.
Both modules store all data on-device with no network communication

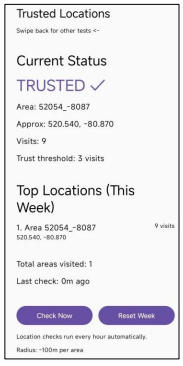

Conclusions

This project shows that a 15-feature behavioral profile built from standard Android sensor data can authenticate users in real time with a mean latency of 72.7 microseconds, around one thousand times faster than the target. The system correctly identifies the enrolled user in 80% of attempts under ideal conditions, all without root access, cloud connectivity or resource-intensive models. The main limitation is Android sandboxing, which restricts gesture monitoring to within the app. Future work should explore system-wide monitoring via the Accessibility Services API and a larger formal participant study to produce statistically robust FAR and FRR measurements.

Acknowledgments

Thank you to Mr. Colin Manning and Mr. Ted Scully for helping me with the project throughout the planning and implementation phase.

An Intrusion Detection System using Machine Learning

Darragh Kelly, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Problem and Aim

Modern networks face increasingly complex cyber threats making traditional signature based Intrusion Detection Systems (IDS) less effective. Traditional systems struggle to detect unknown attacks and often generate high false positive rates (FPR). This project explores the use of machine learning to improve intrusion detection by identifying both known and unknown threats while reducing false positives. The goal is to evaluate different machine learning models and determine which provides the best balance between detection performance and reliability.

Results & Model comparison

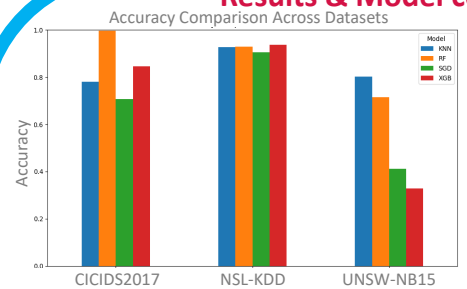

- Random Forest has near perfect accuracy on CICIDS2017
- High NSL-KDD accuracy consistency across all models
- UNSW-NB15 shows a large difference across the models

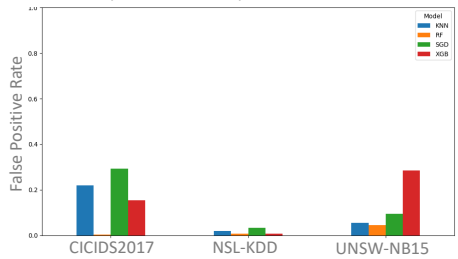

- SGD and KNN have a higher FPR in the CICIDS2017 dataset
- NSL-KDD has the lowest FPR across all models
- Minimising FPR is critical for real-world IDS reliability

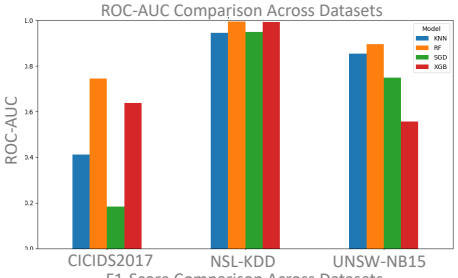

- ROC-AUC highlights how well models differentiate between attacks and normal traffic
- All models perform very well on NSL-KDD

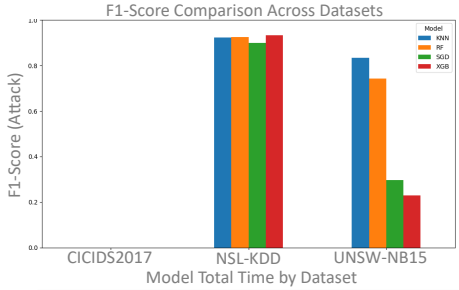

- F1-score provides a better measure than Accuracy for IDS performance
- CICIDS2017 F1-score is zero because the models failed to predict when an attack was malicious

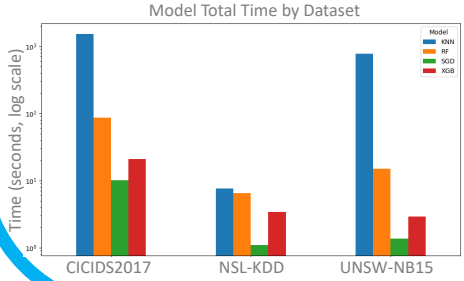

- Time taken is relative to dataset size
- KNN has the longest compute time
- XGBoost and SGD are the most efficient across datasets

Models

This project compares the four machine learning models:

- **KNN:** Simple pattern matching
- **Random Forest:** Advanced non-linear detection
- **SGD:** Fast baseline
- **XGBoost:** Advanced non-linear detection

Datasets

The three datasets representing real world, traditional and modern network traffic are:

- **CICIDS2017:** Real world network traffic and behaviour
- **NSL-KDD:** Traditional Attack types (DoS, Brute Force)
- **UNSW-NB15:** Modern attack types (Exploits, Analysis Attacks)

GOAL: Compare the models and understand which offers the best balance of detection, performance and a low false positive rate.

Conclusions & Future Work

This project showed that machine learning can enhance intrusion detection but also showed that accuracy alone is not a reliable metric due to class imbalance. While the models achieved a high Accuracy score in CICIDS2017 they failed to detect attacks when there were very few attacks for the models to base their predictions on. However Deep Learning approaches using neural networks have the potential to detect complex patterns in network traffic with the ability to improve detection where there are very few attacks.

References

1. G. Kumar, “Evaluation metrics for intrusion detection systems-a study,” Evaluation,vol. 2, no. 11, pp. 11–7, 2014.
2. C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin, “Intrusion detection by machine learning: A review,” expert systems with applications, vol. 36, no. 10, pp. 11 994–12 000, 2009.
3. H. Dong and I. Kotenko, “Cybersecurity in the ai era: analyzing the impact of machine learning on intrusion detection,” Knowledge and Information Systems, pp.1–52, 2025.

Acknowledgments

A special thank you my supervisors Dr Ruairi O’Reilly and Dr Olivia Brickley and my family and friends for their support throughout this project.

Machine Specifications: Intel i7 1255U, 16GB 3200Mhz RAM, 512GB NVMe SSD

ROTEQ - Rowing Statistics Sports Mobile App

Emmanuel Koledoye, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Competitive rowing relies on marginal gains in stroke rate, split time, and power output. Traditional on-water analysis requires expensive specialist hardware, putting elite analytics out of reach for most athletes. ROTEQ addresses this gap by processing smartphone sensor data (accelerometer, gyroscope, GPS) to deliver real-time performance metrics on the water. ROTEQ puts elite-level rowing analytics in the hands of any athlete, no specialist equipment required.

Technologies

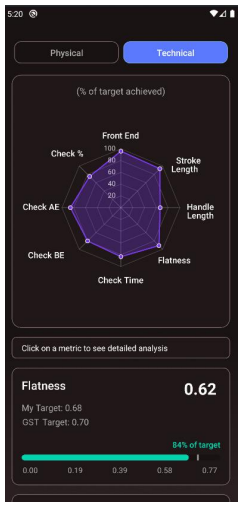

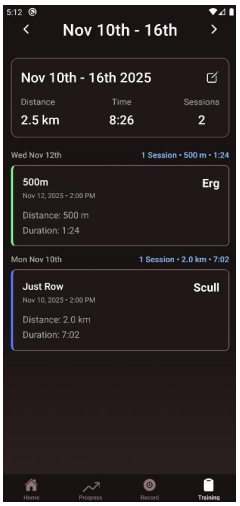

Software: ROTEQ Mobile Application

- Captures stroke data via accelerometer, gyroscope & GPS (iOS/Android)
- Real-time metrics
- Fully customisable screen layout
- Live streaming for remote coach monitoring
- Session history log for trend analysis

Hardware: Weatherproof Mounting System

- Universal mount works on any boat
- Weatherproof enclosure protects from splash & rowing stress
- No interference with the athlete's movement or visibility

High-Level Data Flow Diagram & Architecture

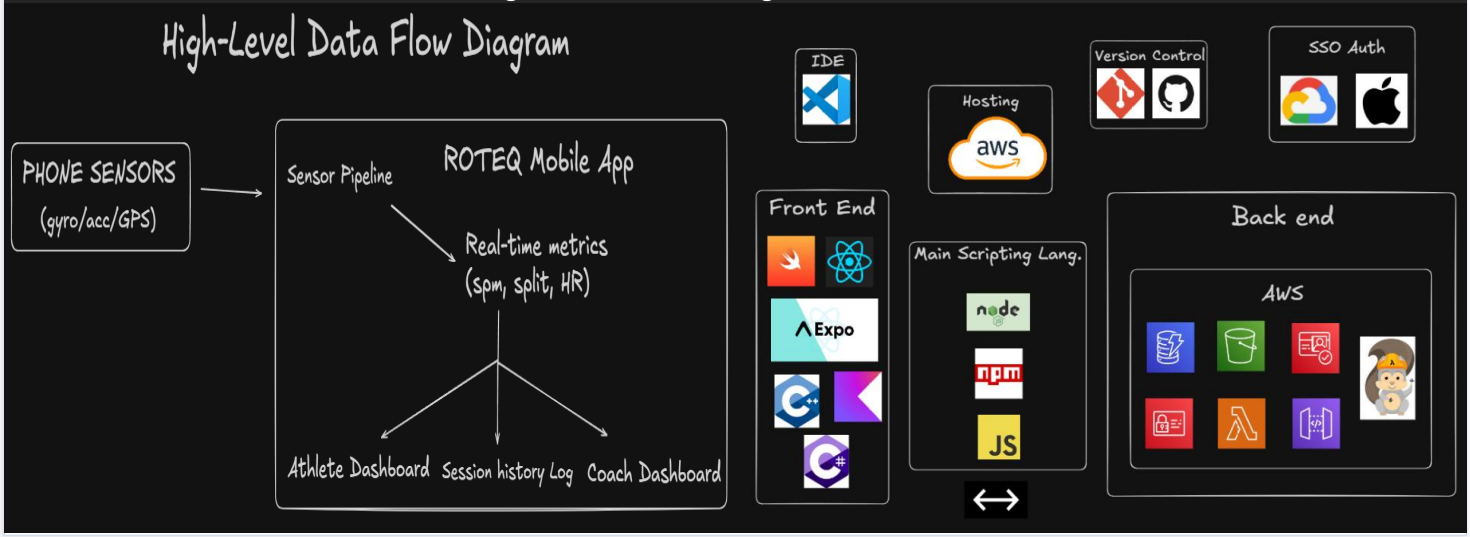

Conclusion

- Smartphone sensors + well-engineered mounting = elite-level analytics without specialist hardware.
- Platform integrates accelerometer, gyroscope & GPS into a clean, actionable interface with live streaming and session history.
- ROTEQ proves that meaningful performance insights don't require specialist equipment, just a smartphone.

Acknowledgments

I would like to sincerely thank the following for their support and guidance:

- Brian Cahill - Project supervisor, Nimbus Gateway Manager
- Brian Murphy - Final year project Co-ordinator
- Conor Mulvey - Nimbus Workplace supervisor, Co-developer
- Will Gunnarsson - ROTEQ Founder, Co-developer
- Sam O' Neill - ROTEQ Founder, Admin

Nimbus Research Centre

Driver Drowsiness Detection Using YOLO Computer Vision Models

Maksymilian Kurzak, BSc Honours in Computer Systems
Department of Computer Science,
MTU Cork, May 2026

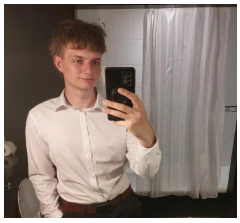

Introduction

The Project's aim is to implement a low-cost plug-and-play embedded device for older vehicles to identifying driver fatigue while driving. With about 20% of deaths related to fatigue-related driving, the project utilizes a Raspberry Pi Zero 2W & Sony IMX500 AI Camera locally for efficiency. It runs without relying on cloud meaning minimal delays and safeguards user privacy. A YOLOv8n model is deployed directly on the IMX500 NPU, enabling for 30 FPS inference on a constrained embedded system. A companion app allows for rapid alert and interaction using Bluetooth Low Energy.

Objectives

- Develop a YOLOv8n with real-time AI inference on IMX500 and embedded hardware
- Deliver real-time alerts with Bluetooth to mobile application to notify driver
- Perform empirical evaluation of varying detection methods (YOLO, DETR Transformer, EAR/MAR Methods)
- Validate performance with real-world tests across varied environmental conditions and participants

Model Evaluation (Training)

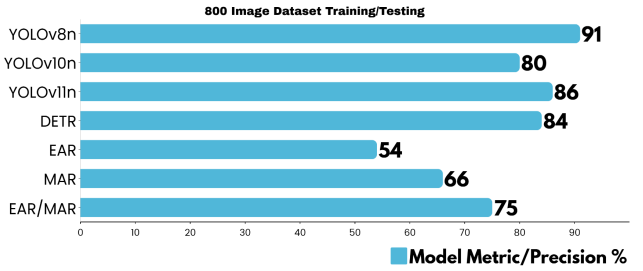

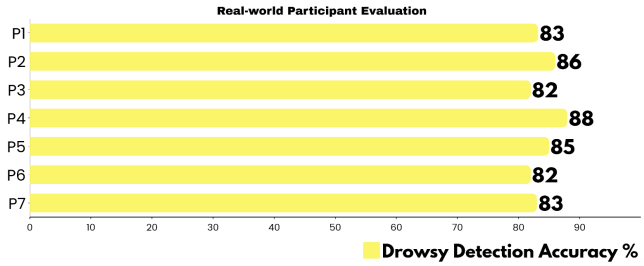

The 1st chart shows YOLOv8n achieving the highest precision, outperforming newer YOLO models, DETR, and geometric methods. This validated v8n as optimal choice for this solution. Bottom chart confirms the model consistency across multiple real-world participants, ranging from 82-88% (85% mean). Controlled 10-minute sessions have been conducted, demonstrating reliability under varying facial characteristics, behaviors, and settings.

System Architecture
Inference to Visualization

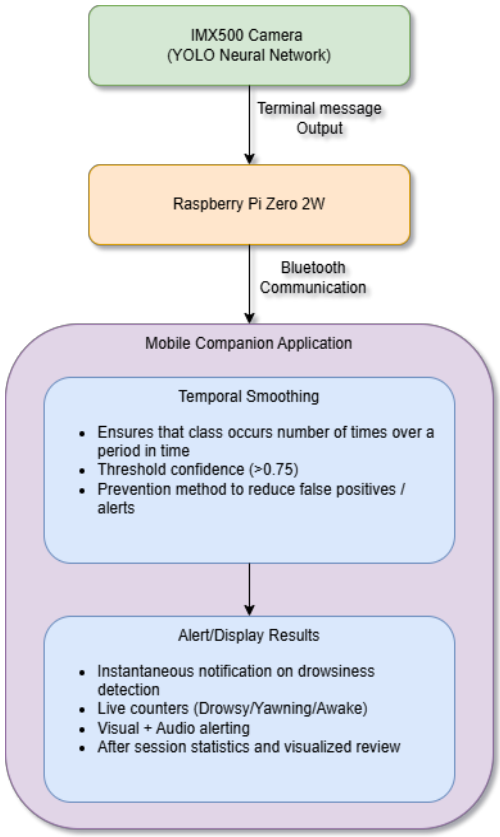

Conclusions

Drowsiness detection with YOLO models showed promising results with real-time performance and accuracy. YOLO outperformed a range of other detection techniques and has shown metrics of 91% precision, 92% mAP50, and 94% recall in many instances. Temporal smoothing has allowed for mitigation of false positives and increasing reliability. In all, an affordable plug-and-play system has been developed, with instantaneous Bluetooth alerting and preserving user privacy. Allowing for closing the gap between affordability and reliable driver monitoring systems.

Acknowledgments: I would like to thank my supervisor Ted Scully for his help and support with this project

Smart SEO and Web Accessibility

Auditing, Testing & Improving the Modern Web

Christopher Luddy, BSc Honours in Computer Systems
Department of Computer Science,
MTU Cork, May 2025

Introduction

- **PROBLEM:** Many websites fail to meet basic SEO and accessibility standards, despite both being essential for a high-quality, inclusive web. Poor accessibility directly excludes users with disabilities (vision, hearing, motor, and cognitive needs).
- **AIM:** Explore how organisations can audit and improve SEO and accessibility using automated and rule-based methods, benchmarked across Amazon, Trinity College Dublin, and Gov.ie.
- **METHOD:** Audits conducted using Google Lighthouse, Screaming Frog, and a custom FYP tool, with results compared to identify gaps and unique findings per site.

Methodology & Tools

Tool	Description
Google Lighthouse	Browser-based automated scoring for SEO, accessibility & performance. Linked to Google search ranking signals.
Screaming Frog	Desktop crawl tool extracting metadata, heading tags & canonicals. Produces no SEO or accessibility scores.
FYP Audit Tool	Purpose-built tool producing SEO & accessibility scores with actionable recommendations based on WCAG 2.1 & Google guidelines.

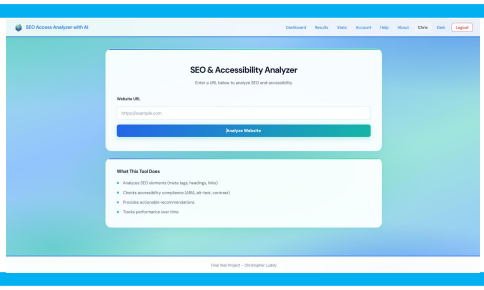

SEO Audit & Findings

Amazon	Trinity College Dublin	Gov.ie
E-commerce	Education	Government
SEO score	SEO score	SEO score
LH: 100	LH: 83	LH: 92
FYP: 25	FYP: 65	FYP: 75
SF: N/A	SF: N/A	SF: N/A
Page load time	Page load time	Page load time
34.3s (LH)	N/A (LH)	9.8s (LH)
4.9s (FYP)	8.25s (FYP)	5.22s (FYP)
5.0s (SF)	6.5s (SF)	4.5s (SF)
Issues flagged	Issues flagged	Issues flagged
ARIA role mismatches	Missing meta description	No meta description (all 3 tools)
Titles <30 chars	Titles <30 chars	Multiple H2 tags
Missing meta descriptions	No canonical URL	Missing anchor text
Missing H1/H2	No Schema.org data	Invalid HTML (SF)
Missing canonical tags	Multiple H2 tags	

SEO Score Comparison by Tool & Sector

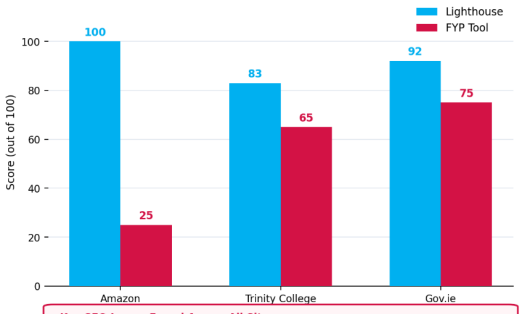

Key SEO Issues Found Across All Sites:

- Missing meta descriptions
- Page titles below 30 characters
- Missing canonical tags
- Missing H1/H2 heading tags
- Missing Schema.org data
- Non-descriptive link text (TCD)

Accessibility Audit & Findings

WCAG 2.1 Compliance

Amazon	Trinity College Dublin	Gov.ie
E-commerce	Education	Government
Accessibility score	Accessibility score	Accessibility score
LH: 92	LH: 93	LH: 100
FYP: 35	FYP: 68	FYP: 62
SF: N/A	SF: N/A	SF: N/A
Violations found	Violations found	Violations found
ARIA attributes not matching element roles	12 images missing alt text	Images using IDs as alt text
Largest gap between LH & FYP across all sites	WCAG 1.1.1 — non-text content	Missed by both LH & SF — FYP only

Accessibility Score Comparison by Tool & Sector

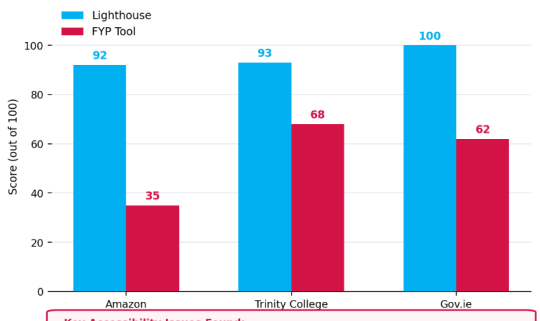

Key Accessibility Issues Found:

- Amazon: ARIA roles not matching element purpose (missed by LH)
- TCD: 12 images missing meaningful alt text — WCAG 1.1.1
- Gov.ie: Images using IDs as alt text — missed by LH (scored 100) & SF

Conclusions

- No single tool provided full coverage; combining all three is necessary.
- The FYP tool caught critical issues Lighthouse missed (e.g., high score but WCAG 1.1.1 violations like incorrect alt text).
- Missing meta descriptions, short titles, and absent canonicals were common, showing high scores do not ensure compliance.
- **These findings highlight the project's value:** linking web audits with accessibility through automated AI scoring for more reliable results and a targeted inclusive web for all.

Issues Detected by Tool Across All Sites

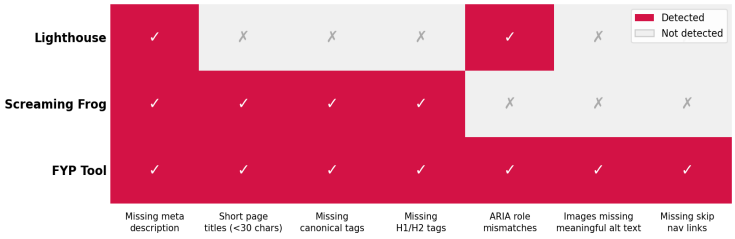

Can We Reduce Energy Waste in Public Buildings Through Machine Learning and Data Analytics?

James Daniel Mallon,
BSc Honours in Software Development
Department of Computer Science,
MTU Cork, April 2026

Background and Objectives

- Governments must reduce emissions due to rising energy costs and EU pressure - Climate Law mandates climate neutrality by 2050 [1]
- Buildings account for ~28% of global carbon emissions [2]
- Cork County Council manages over 500 buildings across Cork County - in 14 of these, smart meters have been installed to monitor heating oil consumption

Problem: Energy waste is difficult to detect using manual analysis

Objective: Develop a system to automatically detect irregular oil consumption and report it to end users in a meaningful way to expedite action

Dataset

- ~5 years of hourly oil consumption data (2020-2025)
- ~346,000 total readings
- Building Metadata (floor area, occupancy, BER)
- Historical Weather Data

3-Stage Hybrid Machine Learning Framework

A hybrid three-stage framework was chosen to combine outlier removal, predictive modelling, and statistical validation for robust anomaly detection.

Dataset Constraints Affecting ML Performance

- Extreme zero-inflation
  - 96.7% of hourly readings = 0
  - 51.6% of daily aggregated readings = 0
- 600+ day gap (June 2023 - March 2025)
- Faulty meters with inconsistent spikes

Machine Learning

Multiple modelling strategies were explored to address performance limitations caused by dataset constraints, with each iteration attempting to improve on the previous.

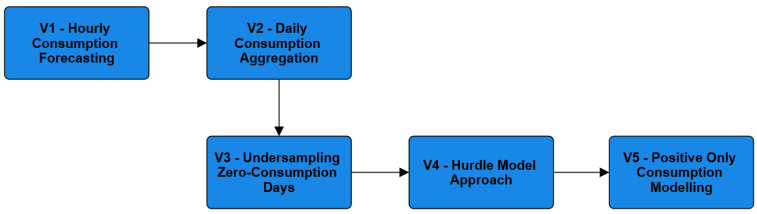

Despite multiple machine learning efforts, performnce remained constrained by the underlying dataset characteristics.

Rule-Based Alerting

- Given the limitations of the dataset, a rule-based approach was adopted to identify irregular consumption patterns based on expected operational behaviour.
- 3 rules were defined and validated by stakeholders:

1. Oil consumption when mean outside temperature > 15°C

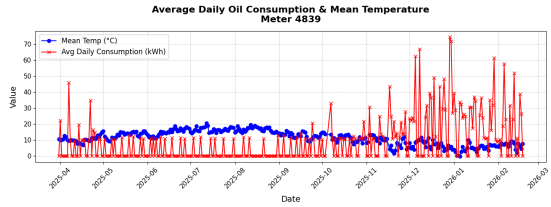

2. Oil consumption outside of normal working hours

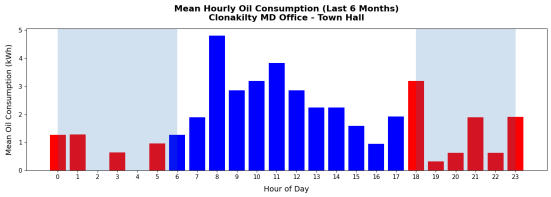

3. Oil consumption on weekends / days when buildings are closed

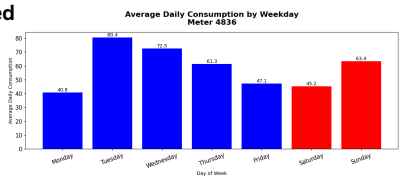

Additional Alerts

- Flag Consumption exceeding a dynamic threshold:
  - mean + 2.5 x standard deviation (per meter)
- Missing Meter Readings:
  - meters with no readings for 24 hours

Alert Implementation

- Framework implemented as a microservice
- Daily alerting script fills alerts table
- Alerts are presented in a user freindly interface on the existing Elighthouse web page

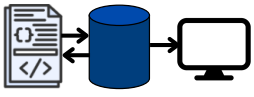

Conclusions

- ML success in literature does not guarantee performance in real-world datasets
- Rule-based system provides immediate, actionable insights in a user freindly way
- A practical data-driven solution has been delivered to CCC

References

[1] (2025) European climate law. [Online]. Available: https://climate.ec.europa.eu/eu-action/european-climate-law en
[2] (2025) Embodied carbon. [Online]. Available: https://worldgbc.org/climate-action/embodied-carbon/

Acknowledgements

The Author would like to thank Dr. Alex Vakaloudis for his support and guidance throughout this project.

Automation Of Penetration Testing Tools

Cormac O Donovan, BSc Honours in IT Management and Cyber Securty Department of Computer Science,
MTU Cork, May 2026

Introduction

The ability for a small business or personal project to withstand a cyber attack can be the difference between success and failure.
anti-virus software checks for malicious file's not for vulnerabilities you may be subject to.
Not only this but anti-virus software may not work all the time. This is why successful cyber attack are becoming an ever growing problem.

Solution

The only way to truly tell if your systems have any vulnerabilities is to preform real world tests. However many people don't have the skill or the knowledge to do this themselves. My proposed solution is to automate real tools used by attackers every day and creating a simplified output that anyone can understand. So i selected three tools to automate in an attempt to strengthen the cyber security space

Nmap

Often one of the first tools an attacker will use

Can reveal exploits through open ports and vulnerable

```
Results Below:
22/tcp  open  ssh         OpenSSH 7.2p2 Ubuntu 4ubuntu2.4 (Ubuntu Linux; protocol 2.0)
80/tcp  open  http        Apache httpd 2.4.18 ((Ubuntu))
139/tcp open  netbios-ssn Samba smbd 4.3.11-Ubuntu (workgroup: WORKGROUP)
```

Attackers then use this information to search for active exploits resulting in access to the system

SQL Map

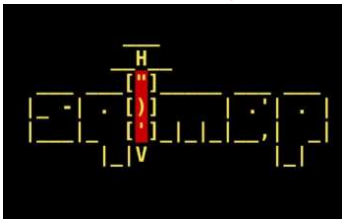

This is use to attack website

It works by attacking the database as most database's are based in SQL

If successful it can result in the entire database (including all user's passwords) being exposed

Example Exploit
User_id=' OR 1= 1;/*
Password=*/--

Query Result
Select*from Users where user_id = '' OR 1= 1;/* ' and password = '*/--'

John the Ripper

Passwords are often encrypted upon storage

Using this tool with a strong password list i can break the encryption exposing the password

If a user has a strong password they become protected against an attack like this

iloveyou
MD5:f25a2fc72690b780b2a14e140ef6a9e0

Conclusions

Automation of these tools could change the cyber security space forever
Anti-virus software only scans for malware but dose not scan for active vulnerability on your network
By following the steps that attackers follow we can understand how to prevent the attacks from working – this is the exact principle of penetration testing

Acknowledgments

The author would like to Ammar Zeini and Larkin Cunningham for being my supervisors for this project

Interactive Data Visualisation Tool for Pokémon Platinum Nuzlockes

Liam Ó Dubhgáin, BSc Honours in Software Development
Department of Computer Science

Introduction

In Pokémon Platinum, the high-stakes 'Nuzlocke' challenge, which is defined by permadeath and restricted encounters, demands precise decision-making in order to successfully complete a playthrough. However, current text-heavy resources hinder rapid data interpretation. This project introduces an interactive visualisation platform to facilitate efficient data comparison and strategic team-building, investigating how visual representations can enhance understanding, engagement, and decision-making within complex, data-rich environments.

Platform UI

Base stats comparison between potential team member (left) and current team members (right)

Type matchup comparison between potential team member (left) and current team members (right)

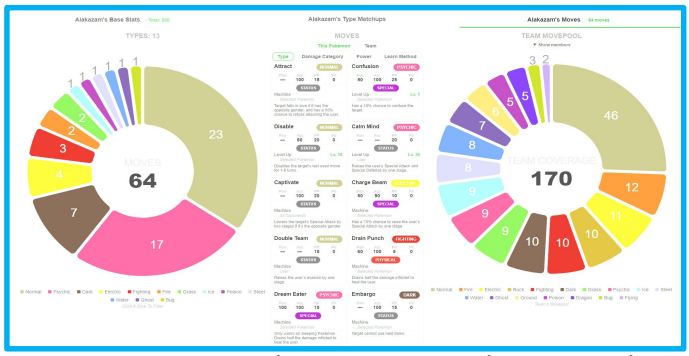

Move set comparison between potential team member (left) and current team members (right)

Results

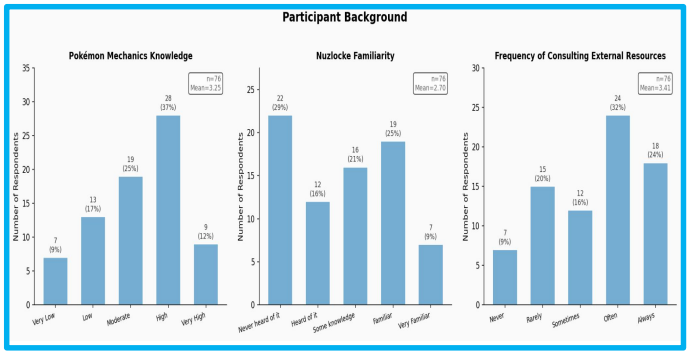

Participant knowledge of mechanics, Nuzlocke challenge familiarity, and external resource use

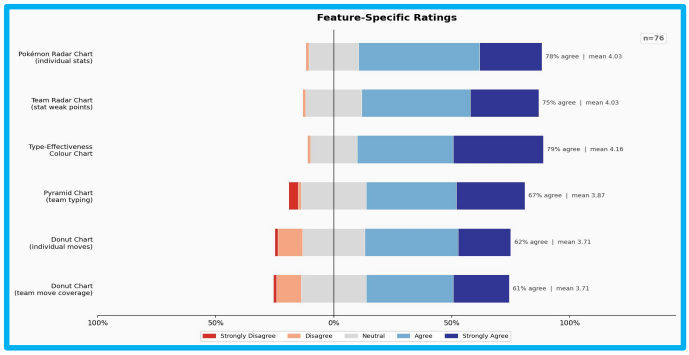

Consensus on the utility and perceived usefulness of the visualisations of key Pokémon characteristics

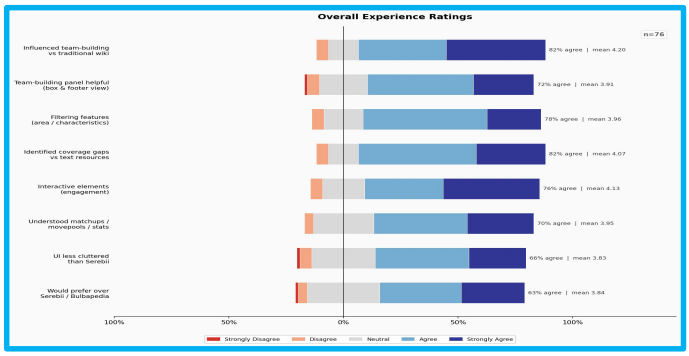

Consensus of the user experience and Quality of Life features of the Platform

Conclusions

Based on the user evaluation of 76 participants, the platform achieved high positive ratings across its visualisations, with participants agreeing they aided in quick comparison of current and potential team members. Users reported improved understanding of complex data and enhanced team building. Most importantly, 63% stated they would prefer to use the platform over existing resources. These findings confirm that integrating interactive data visualisation significantly enhances user understanding, engagement, and decision-making.

Architecture for a Decentralized Smart Device Ecosystem

Chinaza David Ogwudiegwu
BSc Honours in Software Development
Department of Computer Science, MTU Cork, May 2026

Introduction

Traditional smart systems in homes and factories often rely on central hubs, creating single points of failure, added latency, and privacy risks from cloud data transfer.
This FYP Swarm Engine prototype introduces decentralized temperature control using ACO-inspired swarm intelligence, with automatic leader election, distributed peer messaging, and Flask-based web monitoring on simulated Raspberry Pi nodes to deliver resilient, hubless operation.

Methodology & System Architecture

The system uses a three-layer architecture: (1) Device Layer for sensors/actuators, (2) Coordination Layer for distributed swarm agents, and (3) User Layer for dashboard monitoring and manual override. Nodes exchange local state with neighbors and make decisions without a central controller. Development followed an iterative process with modular components for communication, optimization, and visualization. Experiments evaluated convergence behavior, responsiveness, and robustness under node/link disruption.

Architecture

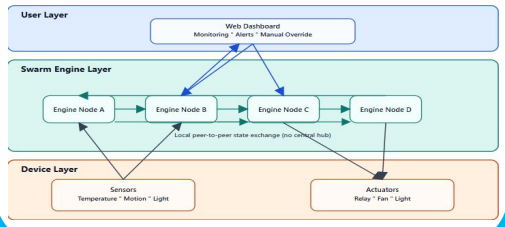

Results

- **Implemented a decentralized swarm runtime with ACO-based control.**
- **Added leader election, heartbeats, and stale-peer handling.**
- **Built LAN messaging with UDP discovery and TCP ACK/retry delivery.**
- **Integrated security: HMAC signing, encryption, and replay checks.**
- **Developed a Flask dashboard with live state and target override API.**
- **Added automated experiment runs with JSON/CSV reporting.**
- **Executed multi-node tests on Raspberry Pi VMs.**

Node map with elected leaders

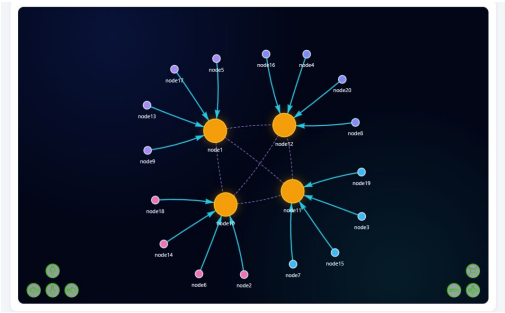

20 nodes

Temperature Nodes Converging

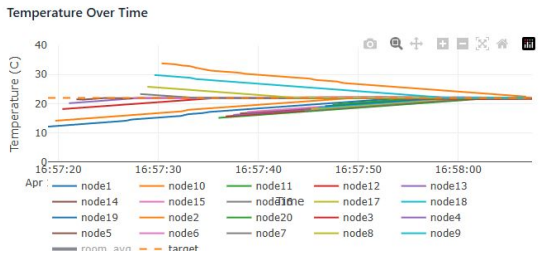

20 nodes

Discussion & Future Work

- **Prototype validates a modular decentralized architecture.**
- **Gateway abstractions exist;**
- **Security primitives are implemented; key management needs hardening.**
- **Scalability framework exists; real Raspberry Pi network tests are next.**
- **Future work: larger fault-injection runs, ACO tuning, baseline comparison.**
- **ZigBee integration is feasible and planned as a future extension for low-power, real-world device communication.**

Conclusions

The architecture is modular and decentralized, but still prototype-level. Gateway abstractions exist, while full Protocol-agnostic integration is pending. Security features are implemented, but key management and hardening are still needed. Scalability tests are automated, but real Raspberry Pi network validation is next. Future work includes hardware deployment, larger fault-injection tests, ACO tuning, and baseline benchmarking.

Technologies Used

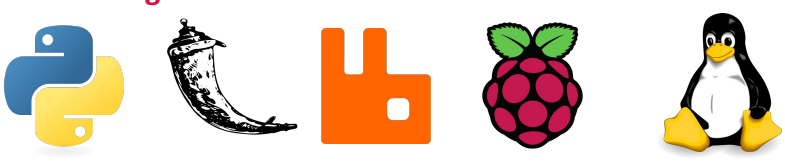

Acknowledgments

The author would like to acknowledge first-semester supervisor Hamdan Awan and second-semester supervisor Meabh O'Connor for their guidance and support throughout this project.

MTU

WATRA

Web Application for Troop Registration and Administration

Wiktor Owczarek, BSc Honours in Software Development, MTU Cork, 2026

Introduction

Scouting is one of the most influential youth movements globally. Millions of participants belong to thousands of organisations, each with their own unique rules and structures.

50,000,000
The estimated number of scouts worldwide [1]

In recent years, scout leaders have been burdened with new responsibilities emerging from regulations and legal requirements. Scout groups have had to find solutions for the collection and management of data of their members and operational paperwork.

The complexity and versatility of scouting structures has made it challenging to implement a universal solution.

The goal

WATRA is a versatile data management platform which aims to ease the burden of administrative processes for any scouting groups, from independent local troops, to some of the largest international organisations.

Highly configurable
Each feature can be adapted to suit diverse needs

Simple and lightweight
A focus on ease of use and performance aims to save time

Open-source and self-hosted
A community driven project, which anyone can host for free

Screenshots

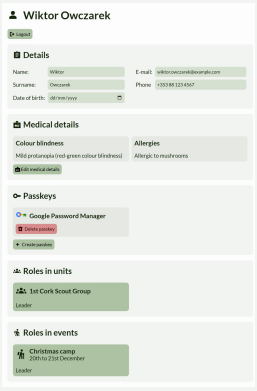

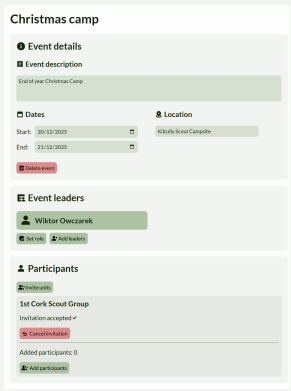

Technologies used

Deno

mongoDB

mongoose
elegant mongodb object modeling for node.js

Eta

Existing solutions

While some groups still use paper records, most use a platform which falls into one of the following categories:

School or sport club platforms
Scouting organisations tend to have more intricate and nuanced structures than the simple student-teacher model offered by these platforms. This mismatch often leads to inefficiencies and hacky workarounds.

Proprietary systems
Platforms developed for a specific large national scouting organisation, kept private from others. Many smaller groups lack resources to develop their own system.

Tipi

Costly SaaS solutions
As non-profits, scouting organisations have limited budgets, making these often inaccessible.

These platforms do not offer the same level of flexibility and customisability as a self-hosted solution.

orgo

Main features

Unit hierarchy
An organisation's hierarchy can be mapped through units and sub-units. New units can be created by local leaders whenever needed.

User management
WATRA elminitates the need for multiple profiles per person. Parent, leader and participant functionality is flexible and determined by role, not account type.

Data collection
WATRA enables the collection of personal details such as date of birth, emergency contacts, medical and dietary information, and user photos.

Data access
User data can be easily accessed, filtered, and exported, by users with appropriate permissions.

Configurability
Every aspect of *WATRA* is data driven, allowing it to be adapted to suit the diverse needs of various organisations. This includes core elements as well as minute details of the system.

Events management
Camps and trips are an integral part of all scouting activities. Events and participant registration can be easily managed by the platform.

Custom data collection and payments
Additional information, consent, and payments can be collected from users through a highly customisable form system.

[1] *WOSM Membership Census 2024*

Dynamic Difficulty Adjustment

In a wave survival game

Rytis Pagalys, BSc Honours in <Computer Systems>
Department of Computer Science,
MTU Cork, May 2026

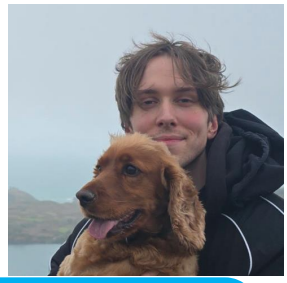

Introduction

Many games use static difficulty systems, which can lead to frustration when the game is either too difficult or boredom when it is too easy.
This project implements a Dynamic Difficulty Adjustment (DDA) system in a wave-based survival game developed in Unity.
The system adapts difficulty in real time based on player performance metrics such as damage taken and health percentage, with the aim of maintaining engagement and providing balanced gameplay.

DDA System

Inputs:
- Damage taken per wave
- Player health percentage (%)
- Wave Performance

Outputs:
- Wave difficulty multiplier
- Enemy damage multiplier
- Health Pickup drop chance

Wave Spawner

- Point budget Based
- Controls enemy spawning using difficulty multiplier.
- Adjusts wave intensity dynamically

Unity matplotlib

Custom-built game developed in Unity

Logging

Logs gameplay data (health, damage, enemies, DDA values)
Records DDA events per wave
Used to generate graphs & evaluate player / system performance

Objectives

- Design and develop a wave-based survival game to provide a controlled environment for implementing and evaluating the DDA system
- Implement a rule based DDA system
- Adapt difficulty based on player performance
- Evaluate effectiveness using gameplay data

Future Work

Future work could extend the DDA system by implementing longer term player performance trends and adapting wider range of parameters such as enemy behaviour, spawn patterns, wave composition. More advanced approaches, such as machine learning could be explored, however, these would require more complex game environments and larger datasets for training.

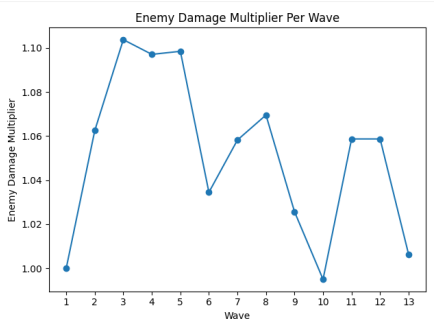

Enemy Damage scales, increasing with strong player performance, decreasing slightly when player struggles

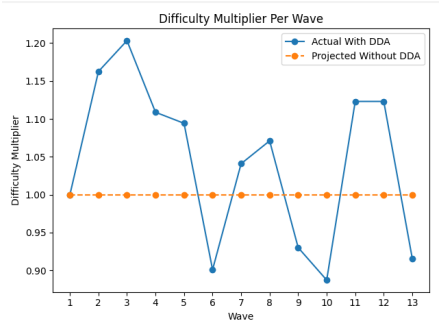

DDA adjusts difficulty above and below the baseline

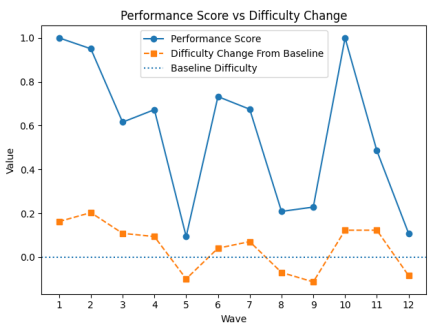

Changes in difficulty reflect player performance

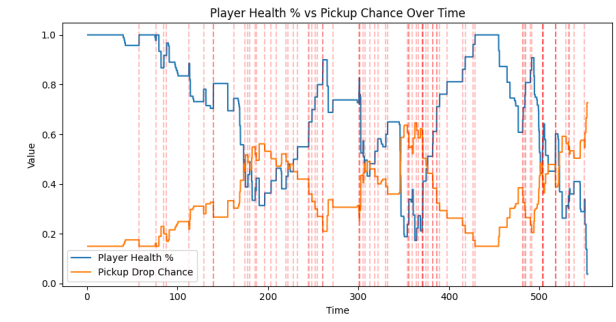

As health decreases, the chance of a health pickup dropping increases, providing support during moments of struggle
*Red lines indicate use of health pickup

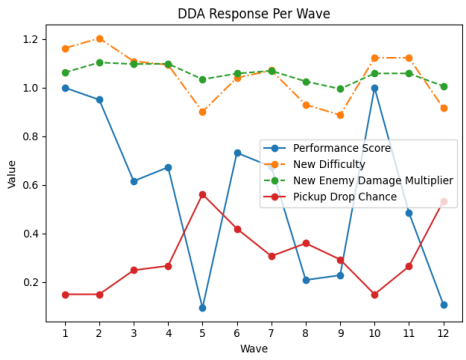

Multiple parameters are adapted in response to performance balancing challenge and assistance across waves

Smoothening (linear interpolation) was used to maintain stability and prevent abrupt difficulty changes

Conclusions

- DDA was successfully implemented in a wave survival game
- Results show that difficulty increased during strong performance and decreased when player struggled demonstrating real time balancing.
- Comparison with a static (non dda) system showed that the DDA approach adjusted difficulty dynamically, whereas a static system would maintain a fixed progression regardless of player performance
- The use of performance metrics enabled responsive adaptation and improved gameplay balance by maintaining level of challenge

ELA: English Language Assistant

An AI-Assisted Framework for Authentic English Comprehension Using Recursive Linguistic Annotation

Vladyslav Rastvorov (R00274535) · BSc Honours in Computer Systems
Supervisor Semester 1 (Research Phase): **Dr. Alex Vakaloudis**
Supervisor Semester 2 (Implementation Phase): **Dr. Nasir Ahmad**
Department of Computer Science · Faculty of Engineering and Science · **Munster Technological University Cork** · May 2025

Problem & Motivation

Mainstream language learning platforms present **pre-simplified content**, stripping the grammatical structures and vocabulary range that define authentic English. Learners are shielded from complexity rather than guided through it. ELA annotates authentic text and speech **in place — structure intact** — using rule-based NLP and fine-tuned ML models. Each sentence is decomposed into a **Recursive Linguistic Element (RLE)** tree enriched with CEFR level, tense–aspect–mood (TAM) tags, IPA, and bilingual notes — processed fully offline, client-side, without cloud dependency.

The RLE Data Model

```
Sentence
  └── Phrase (NP/VP/PP/ADJP)
        └── Word: translation
              CEFR · POS · IPA · TAM · note
```

Language-agnostic · EN→RU via M2M100 (offline).

Analysis Pipeline

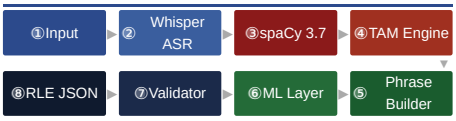

PWA — Android-installable, fully client-side, no cloud.

Research Journey: What We Tried

CEFR CLASSIFICATION

~~DeBERTa~~ underperformed → ~~BERT~~ same → **XGBoost ✓** 78.3%

NOTE GENERATION

~~GPT-4o~~ cost/cloud → ~~T5 raw~~ no generalise → **T5+Templates ✓** ROUGE-L 0.81

ASR

Whisper ✓ WER 8.2%

TRANSLATION

~~GPT-4o~~ opt in → **M2M100 ✓** offline

Context–Pattern Interaction

"In open systems, context forms patterns. In closed systems, patterns form context."

— Discovered empirically during T5 fine-tuning

T5 on raw notation **failed**. Structural templates — **lexically invariant** — were the breakthrough:

❌ `"gave him" → [VP gave][NP him]`
✅ `[V-PAST][IOBJ] → [TAM:past-simple] [VOICE:active]`

The Duality Loop Diagram

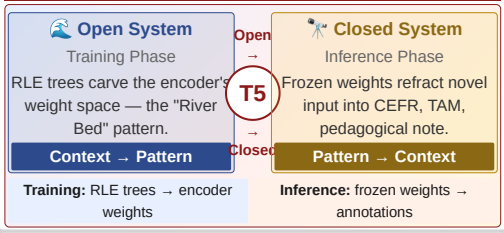

Training Data & Corpus Scale

OANC	30M+ tokens
UD GUM	5.2M tokens
Gutenberg	1M+ tokens
MASC	~500k tokens

27k→685 unique · **6 CEFR** A1→C2

Annotation Backend Comparison

Approach	Quality	Cost	Offline
GPT-4o	High	$$$	✗
Rules only	Medium	Free	✓
T5 fine-tuned	**High**	**Free**	✓
XGBoost CEFR	78.3%	Free	✓

System Validation

- **O1:** spaCy + TAM → schema-valid RLE ✓
- **O2:** CEFR 78.3%; T5 ROUGE-L 0.81 ✓
- **O3:** Visualizer + JSON/CSV export ✓
- **O4:** Android PWA, fully client-side ✓

Performance Metrics

78.3%	0.81	8.2%
CEFR Accuracy	ROUGE-L	ASR WER
XGBoost ensemble	T5 fine-tuned	Whisper base

Conclusions

- AI annotation makes authentic English **accessible without simplification**.
- The **Context↔Pattern duality** generalises to any open→closed system; T5 placeholders are **lexically invariant**.

Future Directions

- **Tauri 2 Desktop App** — native .deb/Windows with bundled ML models.
- **Translation Quality Gates** — ChrF scorer, M2M100 review queue.
- **M2M100 LoRA Fine-Tuning** — domain-specific EN→RU refinement.
- **Real-time Streaming ASR** — Whisper streaming for live practice.

Source Code
github.com/welsol21/FYP_LLM

Live Demo
www.el-a.uk

References

1. Honnibal, M. & Montani, I. (2017). spaCy: Industrial-strength NLP.
2. Raffel, C. et al. (2020). Exploring the Limits of Transfer Learning with T5. *JMLR*.
3. Radford, A. et al. (2023). Robust Speech Recognition via Large-Scale Weak Supervision. *ICML*.
4. Chen, T. & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *ACM KDD*.
5. Levy, M. (1997). *Computer-Assisted Language Learning.* Oxford University Press.

Acknowledgements

The author sincerely thanks **Dr. Nasir Ahmad** (Supervisor Semester 2, Implementation Phase) and **Dr. Alex Vakaloudis** (Supervisor Semester 1, Research Phase) for their invaluable guidance throughout this project. Thanks to the Department of Computer Science at MTU Cork for the academic environment that made this work possible. Inspired by *Eleonora Rastvorova*.

Camera Object Recognition and Search (CObRaS)

MTU
Ollscoil Teicneolaíochta na Mumhan
Munster Technological University

Sofiia Sarana, BSc Honours in Computer Systems
Department of Computer Science,
MTU Cork, May 2026

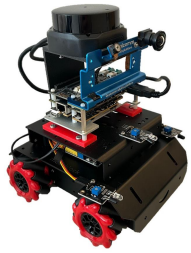

Problem Definition

Object search with mobile robots is challenged by limited onboard resources, large and complex search spaces, high data requirements, and a lack of adaptive learning from experience, necessitating lightweight and dynamic approaches.

- Object Search with Mobile Robots → Limited power and resources
- Large search space → Computationally expensive and long
- Data requirements → Huge datasets, bad generalization
- Search Algorithm Adaptability → No learning from previous experience

Architecture

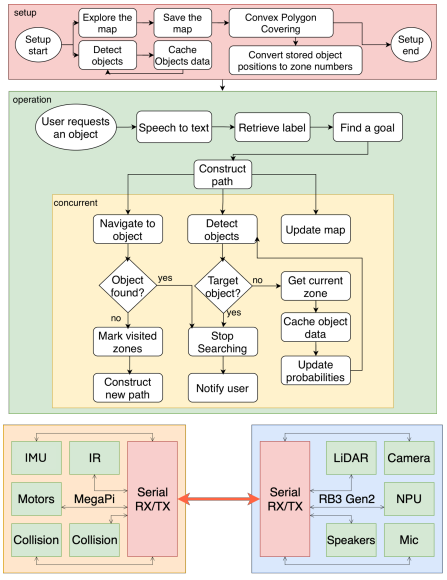

Map & Convex Covering

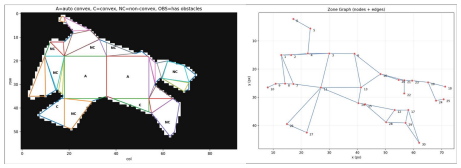

Convex Polygon covering allows the segmentation of the environment into convex zones, which can be evaluated. The graph is constructed based on centroids of each zone.

Raw map → Preprocessing → Grid Generation → Cell Clasification → Convex Decomposition → Size Filtering → Graph Construction

Execution flow

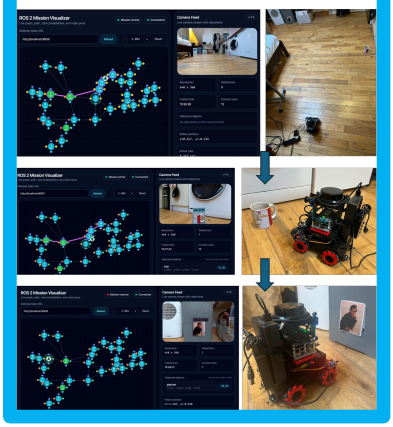

Zone evaluation

$$\log\big(P(O \mid Z)\big) = \log d + \log P_{exact} + \log P_{adjacent} + \log P_{similar} + \log P_{any}$$

$$d = \sum_j a_{o,j} \exp\left(-\frac{\ln 2}{h}(t_{max} - t_{o,j})\right)$$

Time Decay

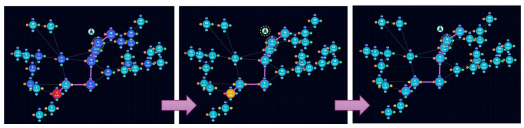

$$P_{\text{exact}} = P(O \mid Z) = \frac{N(z,o) + \lambda}{N(o) + \lambda n}$$

Object being found in a zone

$$P_{\text{similar}} = p(O_{\text{sim}} \mid Z) = \frac{\sum^{k \neq j} s_{jk} p(O_k \mid Z)}{n - 1}$$

Similar object being found in a zone

$$P_{\text{adj}} = P(O \mid Z \in M) = \sum_{z \in M} w_z \cdot \frac{N(z,O) + \lambda}{N(z) + \lambda n}$$

Object being found in adjacent zones

$$P_{\text{any}} = P(Z) = \frac{N_{\text{found}} + k\,p_0}{N_{\text{total}} + k}$$

Any object being found in a zone

Evaluation & Results

Scenario-based evaluation against the baseline algorithm (shortest path to highest probability zone), shows an average improvement of 22%.

Static – best case, fixed object locations
Shuffled – worst case, random location each run
Similar – probability is inferred from similar objects, target never observed
Neighbours – object moves to random neighbour zone at each run
Half – objects are only detected 50% of the time

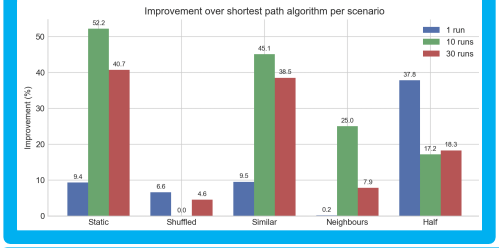

Path Planning

A* with custom heuristic

Maximise the probability along the path instead of relying on immediate goals only

$$g'(n) = g(n) + \alpha \cdot d(n) - (1 - \alpha) \cdot z(n)$$

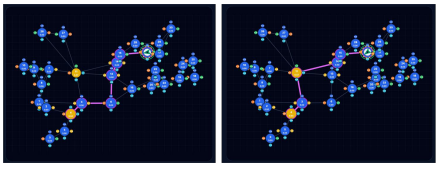

Conclusions

This project presented CObRaS, a zone-based object search system for resource-constrained mobile robots operating in previously unknown environments. By combining online probability updates with probability-aware path planning, the system reduces average search time compared with a shortest-path baseline. Experiments across multiple scenarios showed improvements of up to 40.7% and an overall average improvement of 22%.

Acknowledgments

The author would like to thank Qualcomm for sponsoring this project, and Christian Beder for supervising it .

Camera-less Gaze Tracking

Barnabas Somodi, BSc Honours in Software Development
Department of Computer Science, MTU Cork, May 2026

Introduction

Camera based gaze tracking systems are on the rise in wearable devices, but they have many drawbacks.

Problem Statement

Camera based gaze tracking systems are computationally taxing, relatively large and are expensive.

Objective

The aim is to capture eye motion using low-cost, low-complexity and compact infrared sensors, then compute gaze direction.

Methodology

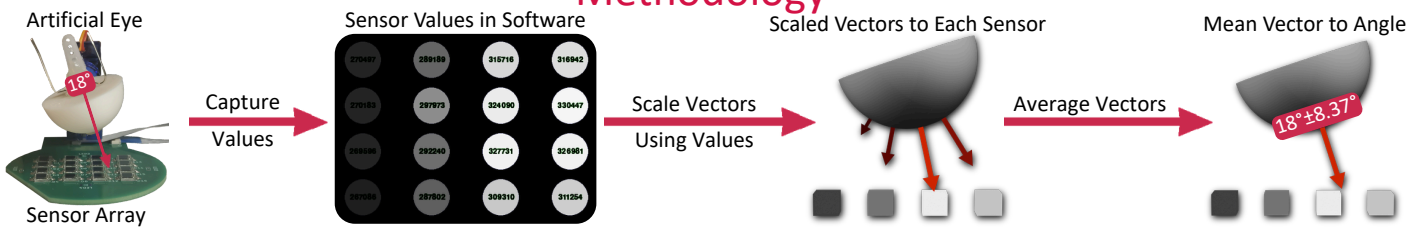

Software & Hardware Architecture

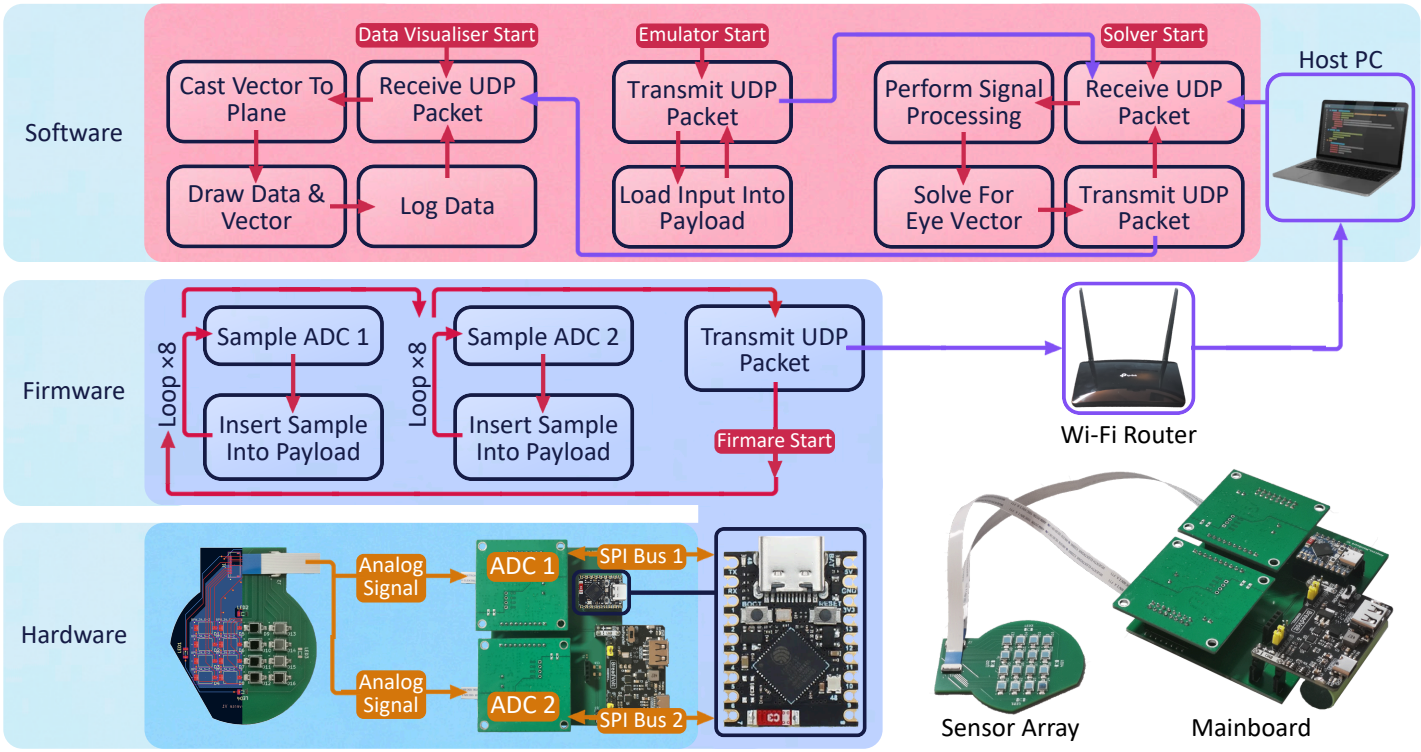

Results

Mean Error by True Angle Bin

	-60° to -50°	-50° to -40°	-40° to -30°	-30° to -20°	-20° to -10°	-10° to 0°	0° to 10°	10° to 20°	20° to 30°	30° to 40°	40° to 50°	50° to 60°
Mean Error (Horizontal)	±17.40°	±14.06°	±10.92°	±7.56°	±6.55°	±5.99°	±5.13°	±4.72°	±1.09°	±4.51°	±10.11°	±15.72°
Mean Erro (Vertical)	±17.46°	±15.99°	±7.38°	±1.94°	±1.07°	±1.70°	±6.34°	±9.18°	±8.05°	±5.60°	±7.35°	±16.61°

-60° (Eye Rotation Limit) — 0° True Angle — +60° (Eye Rotation Limit)

- 120° field of view
- ±8.37° accuracy
- 251hz sample rate

Acknowledgments

Thank you to Dr Hristo Trifonov for supervising this project and Joel Pattison for providing opportunities to practice presenting the project.

Conclusion

While the FOV and sample rate of the system were both much higher than the target, the accuracy failed to meet expectations.

	Target	Result
Field of View	20°	120°
Sample Rate	50hz	251hz
Accuracy	±1°	±8.37°

Quantifying Cross-Modal Interactions in Multimodal Glioma Survival Prediction

From bimodal auditing (InterSHAP) to trimodal integration (MRI)

Iain Swift, BSc Honours in Software Development
Department of Computer Science, MTU Cork, May 2026

Introduction

Gliomas are the most common malignant brain tumours. Even within the same grade, survival varies dramatically due to molecular and cellular heterogeneity. **Multimodal deep learning** combining histopathology (WSI) and genomics (RNA-seq) is widely used for prognosis and performance gains are commonly assumed to arise from ***cross-modal synergy***. But this has never been directly tested in survival prediction.

Do bimodal models truly learn cross-modal synergy, or simply aggregate independent signals?

Does adding MRI as a third modality add measurable prognostic value?

Methodology

Dataset: TCGA-GBMLGG cohort (575 bimodal · 664 with MRI from BraTS2021) · **Two modalities (WSI + RNA-seq) + MRI (FLAIR) as third**

1 · InterSHAP Method

Shapley interaction index adapted from classification to Cox survival. Measures joint-beyond-individual contribution of modality pairs.

$$\varphi_{int} = \tfrac{1}{2}[v(W,R) - v(W) - v(R) + v(\varnothing)]$$

4 model runs per patient · **architecture-agnostic**

Zero → additive · Large → synergy

2 · Fusion Architectures

- **A** **Early Fusion MLP** — 8.8M params · concat
- **B** **Cross-Attention** — 1.8M · Q-K-V across modalities
- **C** **Bilinear Fusion** — 0.54M · FORCES multiplicative
- **D** **Gated Fusion** — 3.2M · dynamic weighting

3 · Trimodal Extension

Add MRI (FLAIR) as third modality via 3D ResNet-18 encoder.

Evaluated across 3 fusion strategies:

- **Early** concatenate features → MLP
- **Late** combine risk scores via Cox
- **Joint** end-to-end encoder training

Controlled comparison: same n=47 train / n=19 test

Results · Bimodal Auditing

Inverse performance–interaction relationship

C-index 0.636 → 0.819 · InterSHAP 4.82% → 3.03%

Better prediction → LESS measured interaction

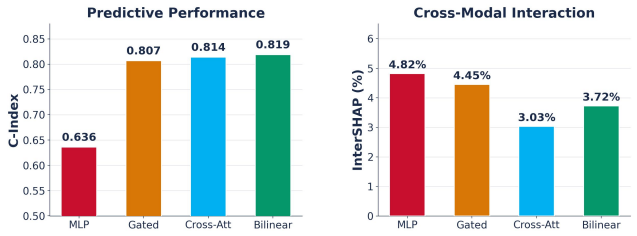

Better models show **LESS** measured cross-modal interaction. Bilinear fusion is mathematically forced to compute multiplicative interactions, yet still shows only 3.72% — lower than the simple MLP.

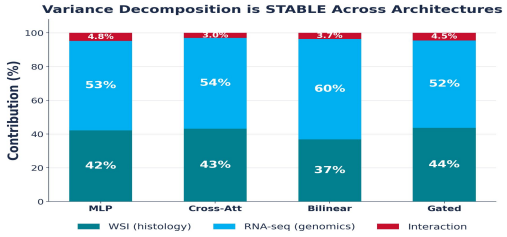

Remarkably stable decomposition: WSI ≈ 40% · RNA ≈ 55% · Interaction ≈ 4%

Across four architectures of wildly different complexity — the split is set by the data, not by the model.

Key quantitative finding

Bootstrap: Cross-Att vs MLP −1.79% (p<0.001) · Bilinear −1.10% (p<0.01)
Synthetic validation: XOR synergy 99.7% · additive <1% (method works)

Results · Trimodal Integration

MRI as a third modality: controlled feasibility test

Trimodal CS = 0.854 · Δ = +0.011 (controlled, n=19)

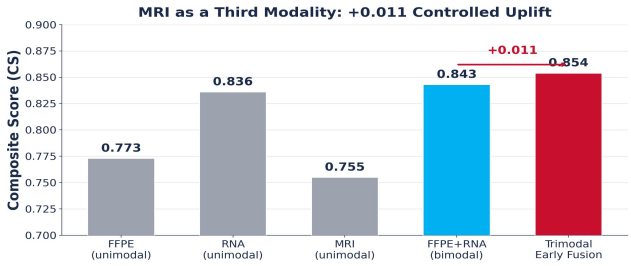

Adding MRI gives a **small but consistent uplift** (+0.011 CS on identical patients). The signal is directional but **not statistically significant** (p = 0.25, permutation test).

Fusion strategy comparison (trimodal)

Strategy	Trimodal CS	Rank
Early Fusion	0.854	★ Best
Late Fusion	0.804	#2
Joint Fusion	0.797	#3

MRI contribution is context-dependent

Neutral-to-negative in bimodal pairs · positive uplift in the trimodal setting
Benefits emerge only with adequate multimodal context

Supporting evidence

MRI unimodal: CS = 0.755 (comparable to histopathology at 0.773)
Bootstrap 95% CI: [0.400, 1.000] — underpowered, n=19 test patients

Conclusions

Measure, Don't Assume

- Performance improvement ≠ interaction learning.
- Cross-modal synergy must be measured directly.
- InterSHAP: 4 evaluations per patient, any model.

Simpler Fusion Works

- Bilinear achieves best C-index (0.819) with 16× fewer params than MLP.
- Complexity improves unimodal encoding — not fusion.
- Invest in better encoders, not fancier mixing.

MRI: Preliminary Uplift

- Trimodal early fusion: CS = 0.854 (+0.011 controlled).
- Direction consistent, but underpowered (n=19).
- Larger MRI-matched cohorts needed to confirm.

References

[1] Wenderoth, L. et al. Measuring cross-modal interactions in multimodal models. AAAI-25, pp. 21501–21509.

[2] Steyaert, S. et al. Multimodal deep learning to predict prognosis in adult and paediatric brain tumors. Communications Medicine 3(1), 44, 2023.

[3] Grabisch, M. & Roubens, M. An axiomatic approach to the concept of interaction among players. Int. J. Game Theory 28(4), 1999.

[4] Mobadersany, P. et al. Predicting cancer outcomes from histology and genomics using CNNs. PNAS 115(13), 2018.

Acknowledgments

Supervisors: I would like to thank my supervisor Dr. Jing Hua Ye for her support and help with this project. I would also like to thank Dr. Ruairí O'Reilly and the Department of Computer Science, MTU Cork.

Paper 1 accepted to XAI 2026 (4th World Conference on eXplainable AI, Fortaleza, Brazil, July 2026).
Paper 2 submitted to IEEE CBMS 2026 (Computer-Based Medical Systems).

Data: TCGA-GBM, TCGA-LGG (The Cancer Imaging Archive) · BraTS 2021 (FLAIR MRI)

Object Detection and Tracking for Pick and Place Robotic Opperations

Ferenc Szabo, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

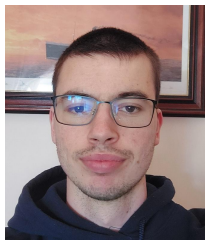

Introduction

This project is meant to investiage the accurancy and rate of collision of a robotic arm when using cameras for picking and placing a moving object.

Challenges

- Achieving fast but accurate object detection
- Fine tuning waist movement during the grasping phase
- Fine tuning object detection as the arm moves closer to the object
- Creating a light weight object detection algorithm

The evalutation setup

System Diagram/Detections

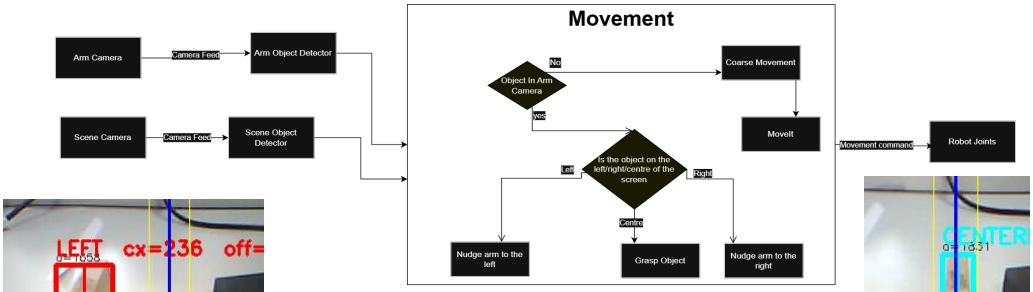

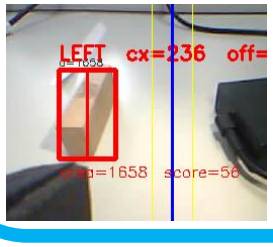

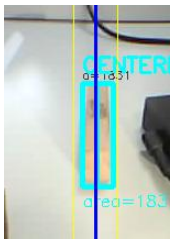

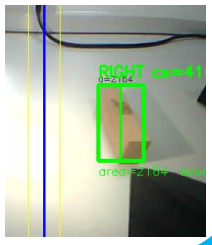

Object Detection

- Object Detection achieved with OpenCv
- Solid object detection based on colour
- ESP32 tracking loop works by sending nudge commands to the waist depending on the location of the object

Movement Planning

- MoveIt library for ROS 2 is used for movement planning
- MoveIt is for big movements and for collision avoidance
- Smaller adjustments for tracking are achieved with joint commands
- Realsense camera uses MoveIt, ESP32 tracking uses joint commands.

Evaluation

- To test grasping accuracy, the number of successful grasps out of 20 was measured
- For a stationary and a moving object
- Lastly the speed was also modified to see how it affects accuracy

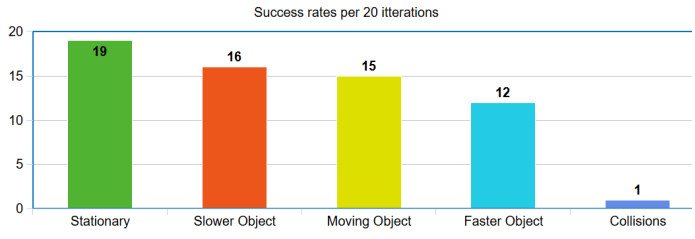

Lessons Learned/Difficulties

- Lighting is really important for an object detection which uses a camera image feed
- Finding the optimal camera placement is really important
- MoveIt is too slow to grasp moving objects.

Acknowledgments

I would like to thank my supervisor ***Victor Cionca*** and my friend Ciaran O Donovan for their support and guidance on this project.

NFC attendance: Tap-In

Kelvin Agu, BSc Honours in IT Managment and Cybersecurity
Department of Computer Science,
MTU Cork, May 2026

Introduction

- The aim of this project to create a system that will aid students and teachers with attendance
- By building and NFC attendance system uses the already available/enable student cards

System Overview

The system uses NFC technology combined with a Raspberry Pi and a web-based backend to automate attendance recording.

Key Components:

- NFC-enabled student cards
- ACR122U NFC Reader
- Raspberry Pi
- Node.js Backend Server
- SQLite Database
- Web dashboard for lectures and students

System Architecture

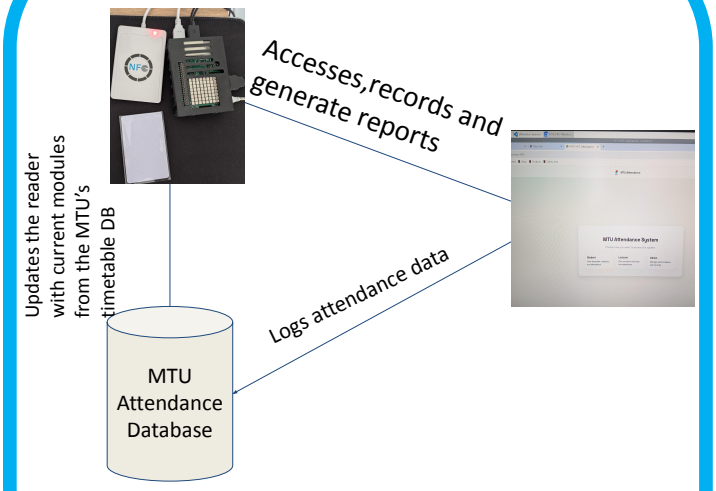

Caption:
System architecture showing data flow from NFC scan to web dashboard

Comparison

Current system	New system
Time-consuming and reduces concentration	Contactless attendance using student ID cards
Prone to human error	Fast check-in process
Allows proxy attendance (students signing in for others)	Automated recording directly into the database
No real-time tracking or visibility	Real-time attendance tracking via web dashboard
Data must be manually entered into systems later	Reduces administrative worklc ırers
Limited verification	

Conclusions

This project demonstrates how NFC technology can be used to modernise attendance tracking in an academic environment. By automating the process through a simple card tap, the system improves efficiency, accuracy, and security compared to traditional methods

Acknowledgments

The author would like to David Stynes and Jing Hua ye.

Real Time Robotic Arm Control via Human Motion Tracking

Stanimir Vasilev, BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

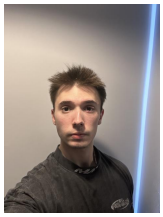

Introduction

- Motivation: enable a robot arm to mimic human arm motion in real time
- Goal: support learning and experimentation in robotic manipulation and human-robot interaction
- Approach: vision-based system that tracks human motion and maps it to robotic control
- Key challenge: human arm has higher degrees of freedom than the robot, limiting motion accuracy and fidelity

System Architecture

The Intel RealSense D435 captures real-time video of arm movement.
YOLOv8 detects 2D keypoints.
Node converts keypoints into joint angles using trigonometry
Angles are sent via ROS2 topics.
The PX150 robot replicates the movement within its limits

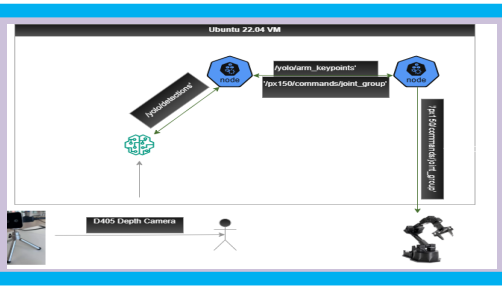

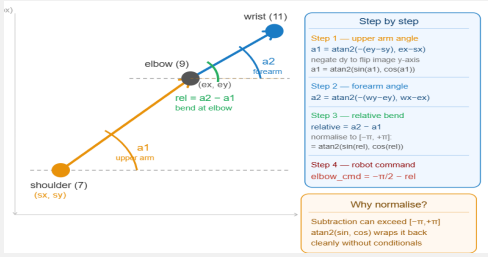

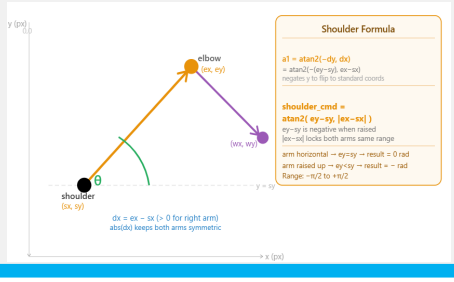

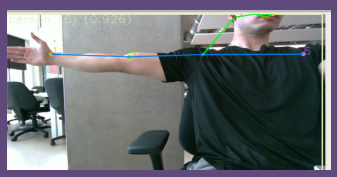
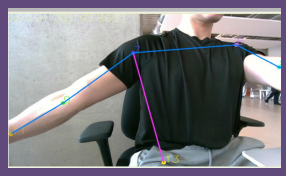
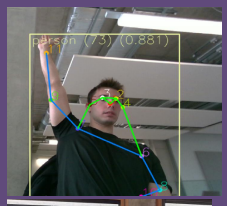
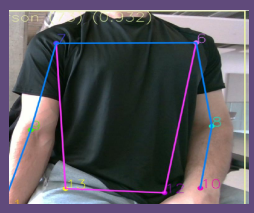



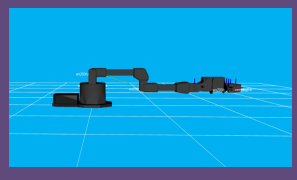

Results

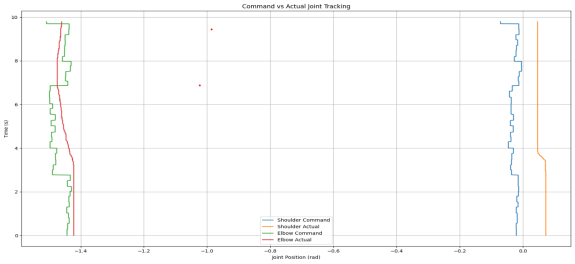

Limitations

- DoF mismatch: human arm is more complex, robot has 4 DoF, accuracy is skewed
- 2D Pose Estimation: Only 2D points provided, no depth included.
- Hardware Limitations: Robot struggles to maintain its weight at certain poses
- Movement constraints: Limited to between -pi/2 to pi/2

Conclusions

- Explored human robot interaction using robotics and AI Built with a student made robotic arm(real systems would be more complex and concise).

References

1. D. Maji, S. Nagori, M. Mathew, and D. Poddar, "Yolo-pose: Enhancing yolo formulti person pose estimation using object keypoint similarity loss," 2022.
2. S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall,"Robot operating system 2: Design, architecture, and uses in the wild,"Science Robotics, vol. 7, no. 66, May 2022.
3. ROS 2 Documentation, *"Understanding ROS 2 Nodes,"*

Acknowledgments

The author would like to acknowledge Victor Cionca for the guidance and help developing this project.

Securing Industrial Networks: A Firewall-Based Segmentation Approach

Isabel Walsh,
BSc Honours in IT Management and Cybersecurity
Department of Computer Science, MTU Cork, May 2026

Introduction

This project examines how network segmentation can improve security in industrial environments, where systems often share the same network. A Cisco ASA 5516-X firewall was implemented to control communication and reduce the risk of unauthorised access between two critical networks.

Objectives

- Design and implement a segmented industrial network architecture
- Enforce least-privilege access between defined network zones
- Prevent unauthorised access and lateral movement
- Align with ISA/IEC 62443 security standards

Proposed Architecture

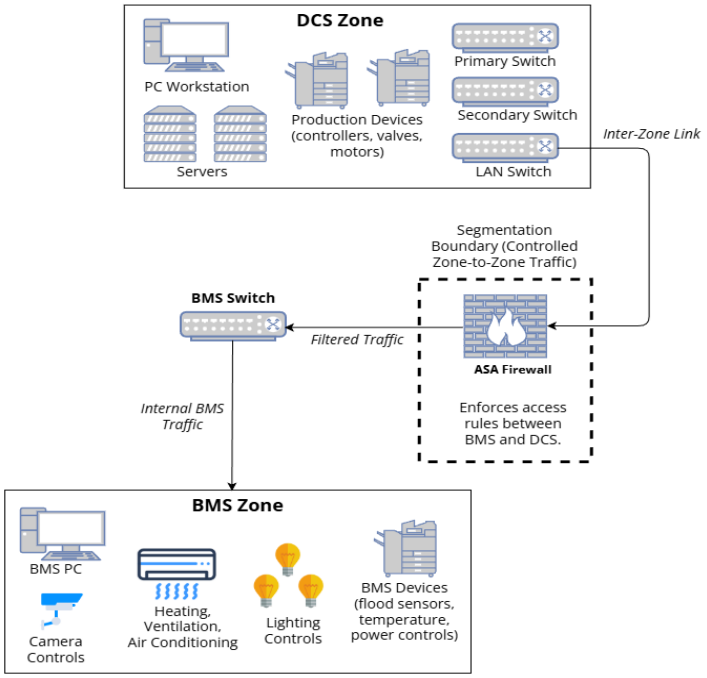

Figure 1: Architecture of the BMS-DCS Segmented Network

Results

Source IP	Destination IP	Description
10.197.250.100	10.197.250.2	Deny icmp src BMS:10.197.250.100 dst DCS:10.197.250.2
10.197.250.100	10.197.250.2	Deny icmp src BMS:10.197.250.100 dst DCS:10.197.250.2
10.197.250.100	10.197.250.2	Deny icmp src BMS:10.197.250.100 dst DCS:10.197.250.2

Figure 2: Unauthorised device blocked from communicating with DCS network

Methodology

Network Setup

- Network devices connected
- IP addresses scheme identified
- Required communication paths and rules identified

Firewall Configuration

- Configured in transparent mode
- Interfaces assigned to network zones
- Network objects and groups made
- Rules/ACLs applied to control traffic

Testing & Verification

- Connectivity validated using ping
- Firewall rules verified via ASDM
- Unauthorised traffic blocked
- Allowed communication verified

Conclusions

- This project demonstrated that firewall-based segmentation can effectively secure industrial networks by isolating systems and enforcing controlled communication.
- The implementation successfully blocked unauthorised traffic while allowing required access, reducing the risk of lateral movement and strengthening overall network security.

Acknowledgments

I would like to thank Jing Hua Ye and Kashif Ahmad for supporting me throughout the project. I would also like to thank the Pfizer Automation Team for giving me the opportunity to undertake this project.

SDOS 'Small DevOps Solution'

DevOps on a standard Desktop Configuration

Ruairi Walsh,
BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Opportunity

As microservice-based architectures open interfaces in key industries, there is an opportunity for small enterprises to accelerate software development and delivery. This project proposes and realises a fully containerised DevOps environment, enabling production-capable workflows locally.

CI/CD Pipeline

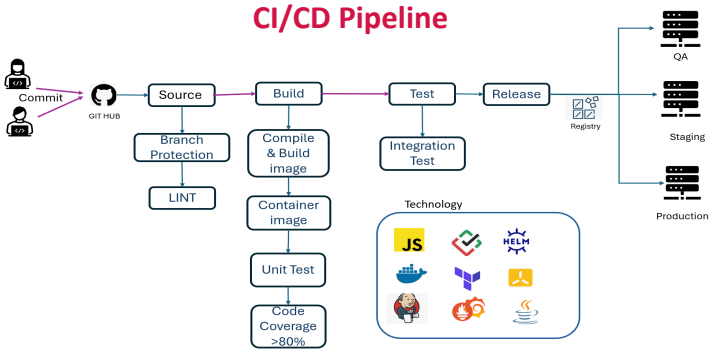

Architecture

Hardware

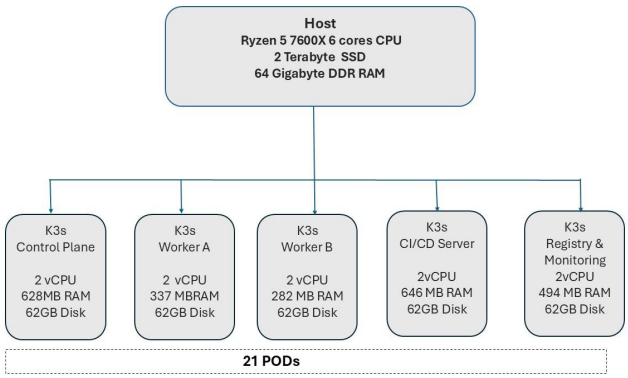

Software

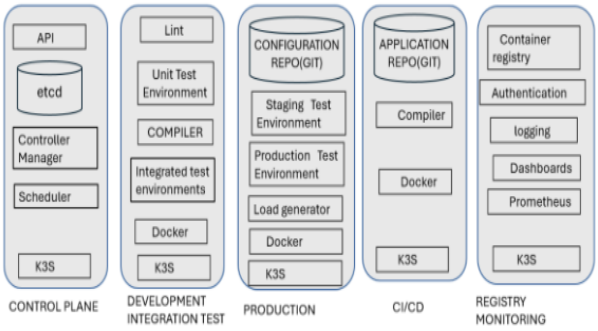

SDOS Dashboards

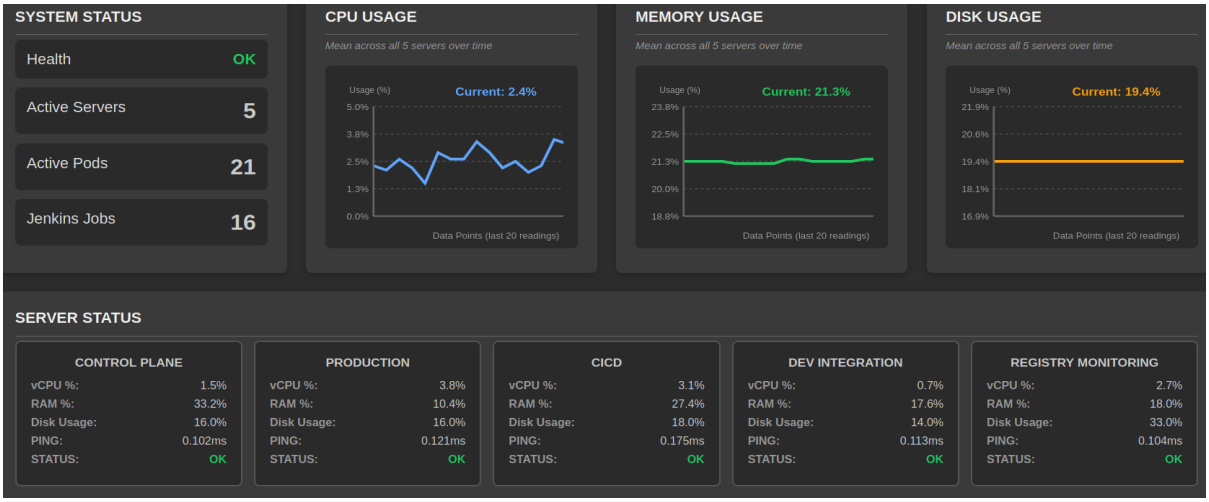

Conclusion

- Using Open source DevOps tools and a standard Desktop platform, this project demonstrated that it was possible to develop and deliver micro-services frequently using a workflow similar to that of a larger organization.

Acknowledgments I would like express my gratitude and thanks to Dr RuairiO'Reilly for his support and help

Pulse Guard

Mobile Integrated Heart Rate Monitoring With Automated Emergency Response

Madison Warner,
BSc Honours in Software Development
Department of Computer Science,
MTU Cork, May 2026

Introduction

Cardiovascular disease is a leading cause of death worldwide, with many incidents occurring outside medical settings

- Many cardiac events happen without immediate medical assistance
- Smartwatches monitor heart rate but lack real-time emergency response
- There is a need for a system that can detect issues and act automatically

Project Overview

This project develops an integrated system for continuous heart-rate monitoring and emergency alerting.

- Smartwatch collects real-time heart-rate data
- Data analysed to detect abnormal patterns
- User is alerted when anomalies are detected
- Emergency contacts notified if no user response
- Combines smartwatch, mobile app, and cloud database

System Architecture

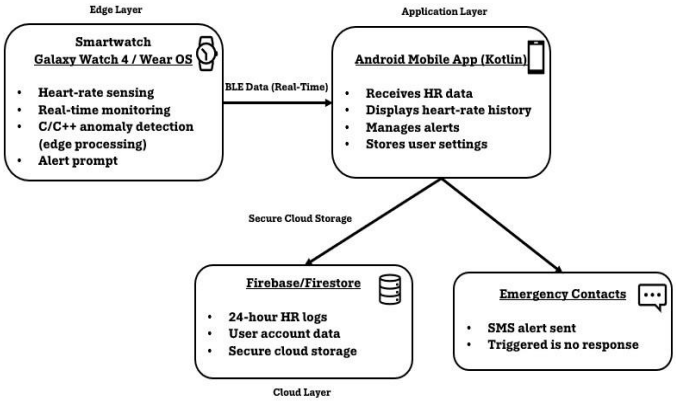

Application

Smart Watch UI

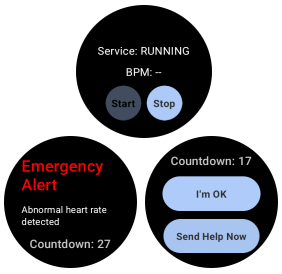

Mobile Phone UI

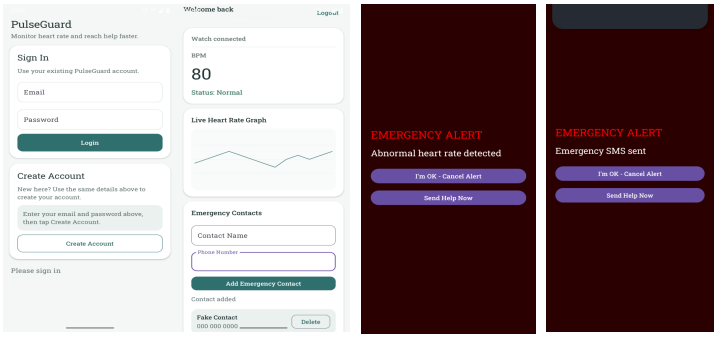

Emergency Message

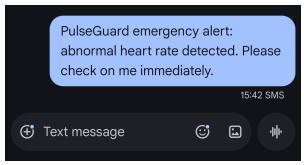

Future Development

- Integrating additional health metrics such as ECG data.
- Machine learning to improve anomaly detection accuracy.
- Direct integration with emergency services and cross-platform compatibility.

Conclusions

By combining continuous heart-rate tracking with automated alerts, the system improves response time in critical situations. It provides a practical, non-diagnostic solution that enhances personal safety using widely available consumer devices.

Acknowledgments

I would like to thank my family and friends for their continued support throughout this project.
Special thanks to my supervisors, Hamdan Awan and Irene Foley for their guidance and feedback.
I also appreciate the support from my lecturers and colleagues who contributed to my learning and development.

Investigating and Mitigating Electro-Magnetic Interference in Semi-Conductor Testing Environments

Brian Wilson, BSc Honours in IT Management and Cybersecurity

Introduction & Motivation

Semiconductor test accuracy depends on repeatable measurements of high-frequency, low-power signals. Modern laboratory environments now coexist with ubiquitous Wi-Fi emissions (2.4 GHz and 5 GHz) from access points and mobile devices.

The EMI Model

Interference in a measurement system can be modelled as an additive superposition of signals:

$$V_{meas}(t) = V_{DUT}(t) + V_{EMI}(t) + n(t)$$

- V_{DUT}: True signal from Device Under Test.
- V_{EMI}: Coupled Wi-Fi interference.
- n: Inherent thermal and system noise.

Coupling Mechanisms

- **Radiated:** Waves traveling through free space received by cables or probes.
- **Conducted:** Propagation along shared power lines or communication buses.
- **Capacitive/Inductive:** Near-field electric or magnetic field transfer between adjacent.

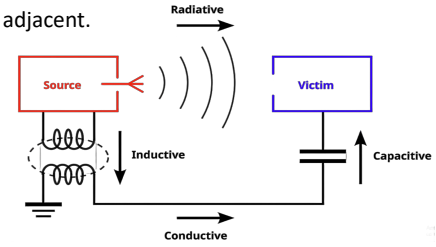

Wi-Fi Activity Heatmap

Top Floor 5GHz

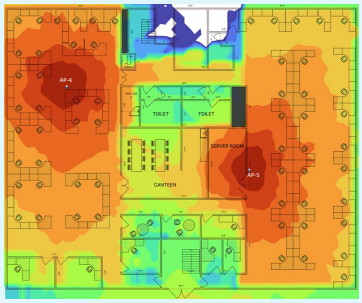

Top Floor 2.4GHz

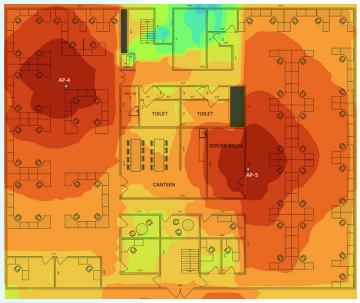

Bottom Floor 5GHz

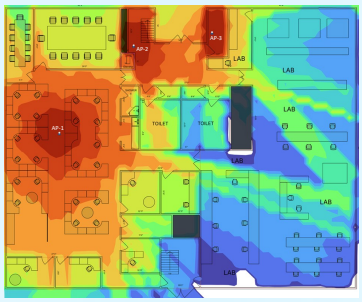

Bottom Floor 2.4GHz

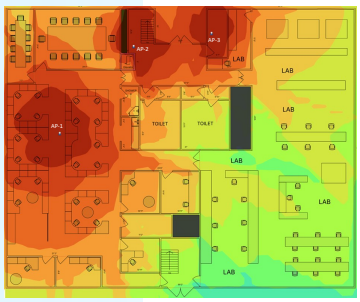

-30 -40 -50 -60 -70 -80 -90 (dB)

Experimental Setup

- **Core Instrument:** A single Vector Network Analyser (VNA) used to capture S-parameters of the Device Under Test (DUT).
- **Representative DUT:** A low noise amplifier (LNA) or RF filter sensitive to external electromagnetic interference.
- **Source of EMI:** A laptop, phone or wi-fi modem

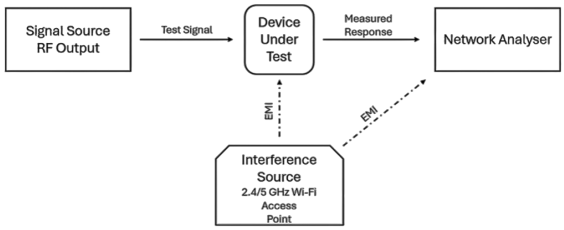

Characterisations & Findings

- **Causality Verification:** Baseline comparisons confirmed spectral spikes are directly linked to router activity and traffic load.
- **Signal Distortion:** Wi-Fi transmissions elevate the noise floor, compromising the repeatability of sensitive RF measurements.
- **Proximity Impact:** Interference intensity correlates strongly with the distance between the Wi-Fi source and test fixtures.

Routers: OFF

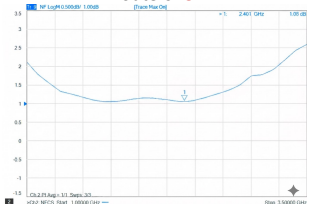

Routers: ON

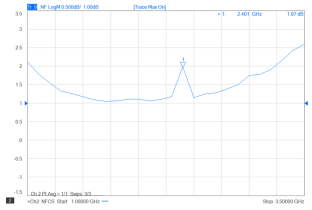

EMI Reduction Results

- **Hardware:** Implementation of Shielded Coaxial Cables and Ferrite Chokes to suppress radiated and common-mode coupling

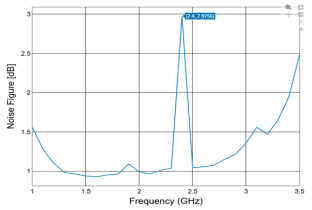

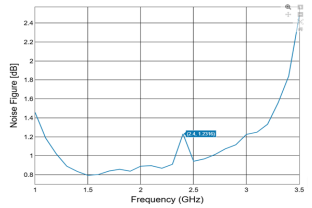

- **Physical:** Use of RF-shielded enclosures to isolate the DUT from ambient laboratory noise

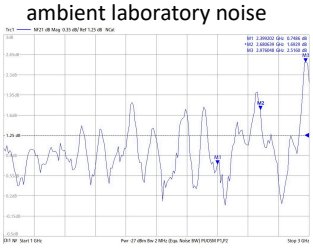

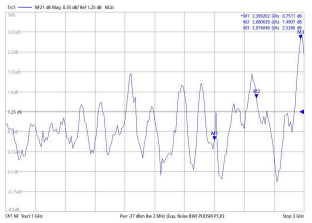

- **Procedural:** Establishing "RF-quiet" zones and scheduling critical tests during low-activity periods

Conclusions

Applying hardware shielding and "RF-quiet" zones effectively suppresses Wi-Fi interference, ensuring the accuracy and repeatability of sensitive semiconductor measurements.

MACOM