



PROCEEDINGS OF THE ANNUAL iCiIT STUDENT RESEARCH SYMPOSIUM **2026**

24TH JUNE 2026

ORGANIZED BY



INFORMATICS
INSTITUTE OF
TECHNOLOGY





**PROCEEDINGS OF THE
iCIIT STUDENT RESEARCH
SYMPOSIUM 2026**

**TRANSFORMING TODAY'S
RESEARCH INTO TOMORROW'S
IMPACT**

24th JUNE 2026

**Informatics Institute of Technology, IIT City Campus, Colombo,
Sri Lanka**

Disclaimer

The responsibility for opinions expressed in articles, in studies and other contributions in this publication rests solely with the respective authors. The iCIIT Student Research Symposium 2026 of Informatics Institute of Technology shall have no liability or responsibility to any person or entity regarding any loss or damage incurred, or alleged to have incurred, directly or indirectly, by the information contained in this book.

Published By:

Informatics Institute of Technology,
435, Galle Road,
Colombo 03
Sri Lanka.

<https://iciit.iit.ac.lk/symposium/>

Published Date: 24th June 2026

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law.

ISSN 3093-5202

© Informatics Institute of Technology, Sri Lanka

iCIIT STUDENT RESEARCH SYMPOSIUM 2026

ORGANIZING COMMITTEE

Symposium Chair

Dr. Ruwan Weerasinghe, Dean Research & Innovation

Organizing Committee

Dr. Dinesh Arunatileka, Head of Research & Innovation

Dr. Thilak Chaminda, Associate Dean Research & Innovation

Ms. Janeesha Shashimini, Research Associate

Ms. Imasha Silva, Research Coordinator

Ms. Hansika Senevirathna, Research Assistant

Editorial Committee

Dr. Dinesh Arunatileka, Head of Research

Dr. Thilak Chaminda, Associate Dean Research & Innovation

Ms. Imasha Silva, Research Coordinator

iCIIT Student Research Symposium 2026 Program

IIT City Campus, Informatics Institute of Technology, 435, Galle Road, Colombo 3

on Wednesday, 24th June 2026 from 9.00 AM onwards

Time	Agenda Item	Details
8:30 AM - 9:00 AM	Registration	
9:00 AM - 9:10 AM	Welcome	Dr. Ruwan Weerasinghe, Dean Research, Computing School, Informatics Institute of Technology
9:10 AM - 9:50 AM	Keynote Speech	Prof. Rangika Halwatura, Senior Professor, Department of Civil Engineering, University of Moratuwa
9:55 AM - 10:30 AM	Refreshments	
10:30 AM - 12:00 PM	Session 1: Presentations	
12:00 PM - 1:15 PM	Lunch break	
1:15 AM - 3:45 PM	Session 2: Presentations	
3:45 PM - 4:15 PM	Distribution of Certificates and tokens of appreciation	
4:15 PM - 4:20 PM	Vote of Thanks	Dr. Thilak Chaminda, Associate Dean, Research, Informatics Institute of Technology
4:20 PM - 4:40 PM	Refreshments	

Table of Contents

iCIIT Student Research Symposium 2026	i
Message from the Chair, iCIIT Student Research Symposium 2026	ii
Message from the Dean, IIT School of Computing	iii
Message from the Dean, Business School.....	iv
ABSTRACTS	1
Track 1: Consumer and Market Insights	2
The Impact of Brand Experience on Customer Satisfaction: the mediating role of Brand Authenticity and Brand Equity in the Hospitality & Beverage Industry	3
Volume to Value: Behavioural Segmentation & Forecasting of Sri Lankan Tourism Markets ..	4
Attitude towards Visual Packaging Design on Purchase Intention: The Mediating role of Brand Experience among Millennials for XXX	5
A Data-Driven Stock Analytics Dashboard for Retail Investors Decision Support.....	6
The Influence of Corporate Social Reputation and Brand Image on Customer Loyalty: The Mediating Role of Customer Satisfaction - A Study on Company A in the Galle District.	7
A Hybrid Ensemble Approach for Detecting AI-Generated Reviews in E-Commerce Platforms	8
Investor 360: An Interactive Financial Analytics Platform based on Benjamin Graham's investment fundamentals for the Colombo Stock Exchange	9
The Impact of Influencer Credibility on Purchase Intention among Millennial Consumers in the XXX Soap Brand through Brand Trust in Colombo District.....	10
A Behavioral Clustering-Based Forecasting Framework for Sri Lankan Kernel-Based Coconut Exports	11
Track 2: Financial, Risk and Healthcare Analytics	12
A Personalised Data-Driven Approach to Predicting Brain Tumour Survivability in Adult Patients in Sri Lanka	13
ForgeDetect: A Hybrid AI-Based Financial Document Forgery Detection System for Credit Lending Verification in Financial Institutions	14
Debtor Tracking and Risk Management Using Machine Learning and Behavioural Analytics	15
SVEMS: AI-Assisted Software Vendor Evaluation Management System for Transparent and Data-Driven Software Selection	16
CeyFleet: Enhancing operational and financial efficiency through effective crew management and skipper recommendation in Sri Lanka's multi-day fishing vessels.....	17

Automated Coco Peat Block Quality and Defect Detection System for Improved Production Efficiency	18
Intelligent Chronic Care Platform - Medikeep	19
Automated Personal Loan Underwriting Using Integrated Rule Based and Behavioral Analytics	20
Brain Tumour MRI Classification Using Ensemble Deep Learning and Explainable AI	21
A Data-Driven Approach to Predicting Hospital Length of Stay, Mortality and Readmission Risk in Myocardial Infarction Patients	22
Track 3: Language Models and Explainable AI	23
Mechanistic Interpretability Under Model Compression for Transformer-Based Language Models.....	24
SINBRIEF: An Unsupervised Framework for Abstractive Text Summarisation of Sinhala Legal Documents	25
CAL-Log: Cost-Aware Active Learning with Logarithmic Cognitive Effort Modeling and Online Adaptation to Human Annotation Behavior.....	26
Apofasi: A Unified Temporally Aware Retrieval-Augmented Generation Framework for Question Answering across Statutory and Case Law in Evolving Mixed Legal Systems.....	27
XFedRec: A Robust, Adaptive, and Explainable Federated Recommendation Framework.....	28
Autodistil-KG: Automated LLM Distillation with Knowledge Graph Supervision.....	29
Quantarra: An XAI Augmented Automated Recommendation Framework for Multi Issue Data Quality Detection in Machine Learning Pipelines.....	30
HEXEval: A Holistic Framework for Standardizing Technical and Human-Centred Explainable AI Evaluation	31
An Interpretable Hybrid Ensemble Framework for Consumer Goods Supply Chain Analytics	32
FGNN: Joint Flow-based Model-level Explainer for Heterogeneous GNNs with Node-Features	33
Track 4: Medical and Health Informatics	34
Hybrid Ensemble of U-Net and Transformer-U-Net for Multi-class Brain Tumor Segmentation	35
Leveraging Knowledge Distillation for Fine-Grained Non-Proliferative Diabetic Retinopathy Classification.....	36
BrainXplain: Towards Explainable Brain Tumor Analysis - Segmentation and Uncertainty Aware Report Generation Pipeline	37

MQML-Inspired Hybrid Quantum-Classical Classification of TCGA-BRCA Tumor and Normal Samples Using Integrated Multi-Omics Features.....	38
Adaptive Label Refinement and Multi-Signal Robust Training for HAM10000 Skin Lesion Classification under Simulated Feature-Dependent Label Noise	39
MindScape: An Explainable Game-Driven Behavioral AI Framework with Graph Neural Networks for Early Cognitive Risk Screening in Children	40
Generative Augmentation of Anatomical-Functional Imaging for Rare Lung Cancer Subtype Classification.....	42
FLEX-Med: A Federated Knowledge Distillation Framework for Acute Lymphoblastic Leukemia Cell Classification under Non-IID Data and Model Capacity Heterogeneity	43
DCTFusion: A Dual Branch CNN-Transformer Framework for Local-Global Feature Fusion in Diabetic Retinopathy Classification	44
MURMUR: Explainable and Personalised Temporal Deep Learning for Breath-Based Augmentative Communication	45
SymAD-ECNN: Geometry Aware Equivariant Learning for Anomaly Detection in Biomedical Imaging	46
Optimizing Adult Glioma Segmentation via Geodesic-Guided Multi-Task Learning and Dynamic Loss Weighting.....	47
Track 5: Robotics, Security and Applied AI	48
AURA: Real-Time Environment Adaptation Using Neuro-Fuzzy Procedural Content Generation with Behavioural Clustering	49
GeMS: Application of Deep Learning Architecture for Multi-Agent Coordination in the Godot Engine	51
A Hybrid Deep Learning Approach for Phishing Website Detection Using Transport and Application Layer Features.....	52
A Mesh-Free, Turbulence-Aware Physics-Informed Neural Network with Explainable AI for Airfoil Optimization at High Reynolds Numbers	53
AIM-IPS: A Five-Layer Adaptive Intrusion Prevention Framework with Cross-Pipeline Threat Amplification	54
Plantar Pressure Heatmap Prediction from Monocular RGB Gait Video using BiLSTM with UNet and Heuristic-based Orthotic Recommendation	55
AeroSynth: A Physics-Informed Multi-Resolution Graph Framework for Aircraft Fuel Prediction and Route Optimization.....	56
Per-Partner OCC Emotional Appraisal and Dual-Process Gating for Sustained Cooperation in Multi-Agent Reinforcement Learning	57

IntelF1: A Machine Learning Approach to Formula 1 Race Winner prediction and Evaluating its Betting Market Efficiency.....	58
DeepEncryptNet: Privacy-Preserving Encrypted Traffic Monitoring via Metadata, Confidence Calibration, and Reinforcement Learning.....	59
Proximal Policy Optimization with Weighted Priority Architectures for Resilient Spatial Navigation and Autonomous Disposal Routing in Mobile Edge Robotics.....	60

iCIIT Student Research Symposium 2026

Transforming Today's Research into Tomorrow's Impact

The iCIIT Student Research Symposium 2026 marks the second edition of this academic platform dedicated to showcasing undergraduate research and strengthening early engagement with scholarly inquiry within the IIT academic community. Building on the success of its inaugural edition, the symposium continues to provide a structured environment where students present their work, receive academic feedback, and experience formal research dissemination.

This year's symposium is guided by the theme "Transforming Today's Research into Tomorrow's Impact", highlighting the importance of translating academic inquiry into meaningful, real-world solutions. The theme reflects the role of student research in addressing contemporary challenges while contributing to future innovation across technology, business, and society.

As part of the broader iCIIT initiative, the symposium reinforces IIT's commitment to fostering a strong research culture that connects classroom learning with real-world application. It encourages students to engage in inquiry-driven learning, critical thinking, and innovative problem-solving, while gaining exposure to key aspects of the research process, including abstract development, peer evaluation, and formal presentation.

The 2026 symposium brings together student researchers from the School of Computing and the IIT Business School, creating a diverse and interdisciplinary environment. This cross-disciplinary participation enriches discussion and promotes collaboration between technological and business perspectives in addressing real-world challenges.

A total of 53 abstracts, presented across five thematic tracks, were showcased this year., covering a wide range of topics including emerging technologies, applied computing, intelligent systems, data-driven decision-making, business innovation, and solutions to contemporary organizational and societal challenges. This diversity reflects strong student engagement and effective research mentorship across both schools.

The symposium features a keynote address by Prof. Rangika Halwatura from the Department of Civil Engineering, University of Moratuwa. The keynote provides valuable insights into current trends and future directions in technology, business, and research, while linking student work to broader global innovation contexts and emphasizing the importance of rigorous and impactful research.

The program also includes interactive sessions that foster dialogue between students, academics, and industry professionals, enabling constructive feedback, idea refinement, and collaborative thinking.

Overall, the iCIIT Student Research Symposium 2026 stands as a significant milestone in the academic journey of IIT students. As the second edition of the initiative, it strengthens the institution's culture of research, innovation, and academic excellence, aligned with its guiding theme of transforming today's research into tomorrow's impact.

Message from the Chair, iCIIT Student Research Symposium 2026



Dr. Ruwan Weerasinghe

It is a privilege to welcome you to the iCIIT Student Research Symposium, a forum that celebrates the very best of undergraduate final-year research undertaken at IIT. This event brings together exceptional student work from the School of Computing and the Business School, made possible through the commitment, guidance, and mentorship of our academic supervisors. The quality and diversity of the research presented in these proceedings reflect not only our students' capabilities, but also IIT's culture of inquiry, innovation, and academic excellence.

Many of the projects showcased here demonstrate the depth, originality, and methodological rigor expected of work suitable for leading international conferences and journals. I strongly encourage students and their supervisors to build on these outcomes and pursue publication in high-quality, Tier 1 venues. Such dissemination gives deserving research the visibility it merits, contributes to global scholarly conversations, and strengthens IIT's standing as a research-active institution.

Equally, research should create value beyond the academic setting. Where a project offers a viable solution to a real-world challenge, I encourage its creators to explore pathways to implementation - whether through the organizations they join, collaborations with industry and community partners, or ventures of their own. May this Symposium serve not only as a recognition of achievement, but also as a starting point for research that is seen, shared, applied, and made meaningful for society.

Dr. Ruwan Weerasinghe

Symposium Chair | Dean Research

Research Department

Informatics Institute of Technology, Sri Lanka

Message from the Dean, IIT School of Computing



Dr. Dilhari Attygalle

It is with great pride that I congratulate the students, supervisors, organisers, and reviewers who have made this year's iCIIT Student Research Symposium a success. The breadth of over 60 paper presentations across three parallel tracks reflects the intellectual curiosity, technical rigour, and determination of our final-year students in the School of Computing.

Research achieves its greatest value when it moves beyond theory to address real needs. This year's focus on converting research into real-world relevance is especially meaningful in computing, where innovation has the power to transform industries, communities, and everyday lives. I commend our students for engaging with complex problems and for demonstrating the potential to turn their ideas into practical, responsible, and impactful solutions.

Dr. Dilhari Attygalle

Dean

IIT School of Computing

Informatics Institute of Technology, Sri Lanka

Message from the Dean, Business School



Prof. P. S. M. Gunaratne

The iCIIT Student Research Symposium is an exciting celebration of creativity, analytical rigor, and professional maturity of our graduating students. Through more than 60 presentations conducted across three parallel tracks, the Symposium has provided a dynamic platform for students to communicate evidence-based insights and engage with issues that matter to organisations, markets, and society.

The theme of translating research into real-world relevance captures a central purpose of business education: generating knowledge that informs better decisions and creates meaningful value. I extend my warmest congratulations to the students from the Business School who contributed to this significant event, their supervisors, and everyone who enriched this year's symposium through their commitment to rigorous inquiry and their readiness to apply research in ways that make a positive difference.

My sincere appreciation goes to the staff of the research unit at IIT, including the Dean of Research, for their commendable work organizing the annual symposium. By creating a leading platform for our students to showcase their finest work to a distinguished audience representing both industry and academia, you continue to bridge the gap between research and practice.

Professor PSM Gunaratne

Dean

IIT School of Business

Informatics Institute of Technology

ABSTRACTS

Track 1: Consumer and Market Insights

The Impact of Brand Experience on Customer Satisfaction: the mediating role of Brand Authenticity and Brand Equity in the Hospitality & Beverage Industry

¹Masharik, A.N.*, ²Ratnayake, G.S.

ainul.20220750@iit.ac.lk, University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka.

Brand experience is often identified as a critical factor that influences customer satisfaction in the hospitality and beverage sector as creating memorable brand experiences lead to building long term relationships with customers. This research study was undertaken to identify the impact of brand experience, brand authenticity, and brand equity on the customer satisfaction level in relation to XXX Tea Lounge.

A quantitative research design was employed and data was collected through a structured questionnaire and SPSS was used to conduct reliability, correlation and regression analysis. Sobels test was used to test for mediating effect. The findings indicate that brand experience and brand authenticity had a significant and positive impact on customer satisfaction..Brand equity, on the other hand, had minimal significance in predicting customer satisfaction. In the proposed conceptual model, brand authenticity was found to be the most significant predictor among the three variables examined. This clearly shows that customers greatly value brands that are considered authentic, trustworthy, and reliable in terms of fulfilling their promises. Brand authenticity and brand equity mediated the relationship between brand experience and customer satisfaction.

Customers are satisfied with brands not only because of their experience with the brand, but also due to increased authenticity of the brands. The findings provide practical implications to brand managers, strategists and marketing consultants on improving consumer satisfaction through improvements in brand experience and brand authenticity. Providing distinctive and authentic beverage experiences and store atmosphere can enhance customer perceptions of brand value and improve customer satisfaction.

Keywords: Brand Authenticity, Brand Equity, Brand Experience, Customer Satisfaction, Hospitality & Beverage Industry

Volume to Value: Behavioural Segmentation & Forecasting of Sri Lankan Tourism Markets

¹Sabry, F.*, ²Kirupananda, A.

mfaadilsabry@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The Sri Lankan tourism sector heavily relies on macroeconomic forecasting for strategic planning, infrastructure development, and budget allocation. However, the traditional predictive approach treats all inbound tourists as a homogeneous entity, relying purely on total aggregate arrival volume. This homogenisation creates a critical analytical blind spot, masking distinct behavioural profiles and resulting in suboptimal commercial investments. This research addresses this limitation by shifting the analytical focus from sheer headcount to economic value, utilising 'Total Bed Nights' as a measurable proxy for visitor commitment.

Following the CRISP-DM framework, an end-to-end analytical pipeline was developed. Fifteen years of historical reports from the Sri Lanka Tourism Development Authority (SLTDA) were extracted, validated, and engineered into specific behavioural metrics, with structural anomalies from the COVID-19 pandemic mathematically mitigated. The methodology employed unsupervised machine learning (K-Means) to segment global source markets into distinct personas. Subsequently, a competitive time-series 'Forecasting Tournament' was executed, evaluating SARIMA, Holt-Winters, and Prophet algorithms, to identify the optimal predictive model for each distinct segment over a 24-month horizon.

The algorithm successfully partitioned the market into four strategic clusters: Traditional Leisure, Resilient Core, Seasonal Youth, and Emerging Markets. The predictive evaluation proved that the combined disaggregated forecasting strategy achieved a highly accurate Weighted Mean Absolute Percentage Error (WMAPE) of 6.80%, drastically outperforming the classical aggregate baseline's error rate of 10.23%.

In conclusion, this research mathematically proves that forecasting distinct behavioural personas independently neutralizes conflicting seasonal noise, resulting in statistically superior predictive accuracy. To bridge the gap between data analytics and executive decision-making, the outputs were deployed into an interactive Power BI dashboard. This provides policymakers with a highly actionable, visual tool to optimise marketing investments and maximize the industry's future economic yield.

Keywords: Behavioural Segmentation, Business Intelligence, K-Means Clustering, Time-Series Forecasting, Tourism Analytics.

Attitude towards Visual Packaging Design on Purchase Intention: The Mediating role of Brand Experience among Millennials for XXX

¹Fernando, A.N.*, ²Rathnayaka, G.S.,
amandi.20221373@iit.ac.lk, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

This study examines how attitude towards visual packaging design influences purchase intention, with brand experience acting as a mediating factor on this relationship, among the millennial consumers for product XXX in Colombo District. The study relies on the theoretical insights including the Stimulus-Organism-Response (S-O-R) model, Visual Design Elements Theory, Cue Utilization Theory, Brand Experience Theory, and the Theory of Planned Behaviour to determine the influence of visual packaging elements on consumer perception and behavioural consequences.

Quantitative research design was used and 396 respondents were sampled using a structured questionnaire. The statistical analysis was done using SPSS, Pearson correlation, multiple regression and mediation analysis. The findings reveal that a positive attitude towards visual packaging design significantly enhances purchases intentions among millennials. Attractive Visual packaging designs positively influence brand experience and lead to improved customer purchase intentions. Mediation analysis establishes that brand experience plays an important role in mediating the relationship between visual packaging design and purchase intention, suggesting that visually engaged packaging creates positive brand experiences as well as improving customer purchase intentions. These findings are important for marketers, branding specialists, product designers and strategic managers as they focus on creating packaging elements that resonate with millennials and create a resonating branding experience.

Keywords: Brand Experience, Millennials Attitude, Purchase Intention, Visual Packaging Design, Wellness industry

A Data-Driven Stock Analytics Dashboard for Retail Investors Decision Support

¹Dias Bellana, S.*, ²Kirupananda, A.
sethmikabellana@gmail.com, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

Retail participation in US equity markets has grown considerably over the past decade, however the tools available to retail investors remain fragmented and underdeveloped. Company fundamentals, price history, valuation metrics, and news are spread across separate platforms in inconsistent formats, while professional systems such as the Bloomberg Terminal are priced beyond the reach of most individuals. This creates a situation where there is plenty of data available but turning that data into useful insights remains difficult. This project addresses that gap through a web-based stock analytics dashboard that brings technical, fundamental, and news-based information together in a single, user-friendly interface.

The system was built using Python and Streamlit and consists of four analytical layers. The technical layer adapts a three-stage signal pipeline combining a 20/50-day SMA crossover, average traded volume slope confirmation, and an RSI filter. A reinforcement learning model using Proximal Policy Optimisation, trained in a custom Gymnasium environment, provides a complementary data-driven layer. A Piotroski F-Score layer covers company fundamentals, while a keyword-based model classifies recent news as positive, negative, or neutral. These components feed into a confidence-based recommendation system supported by plain-English explanations of every output.

The system was evaluated on ten US large-cap stocks between January 2023 and April 2026, generating 8,250 decision points. The rule-based model produced only 34 (0.4%) actionable signals, which suggests that such opportunities are scarce in US large-cap. The reinforcement learning model generated more signals but favoured Buy decisions, so the hybrid system sits between the two within a confidence framework. For the highest-confidence signals, average returns rose from 2.3% after one day to 5.9% after twenty-one days, with win rates improving over the same horizon. Overall, the system achieved 69% accuracy and a Sharpe ratio of 1.06, performing better than random selection across all time horizons considered.

Thereby, two conclusions follow. First, the value of the system lies as much in how clearly it explains its outputs as in raw predictive accuracy, which is a quality that matters greatly for retail users. Second, the scarcity of high-confidence signals suggests that a notification-based approach may serve users better than a dashboard requiring constant monitoring. Remaining limitations, including gaps in Yahoo Finance data, the simplicity of the sentiment model, and the limited extent of hyperparameter tuning, providing clear directions for future work.

Keywords: explainable analytics, reinforcement learning, retail investors, stock market, technical analysis

The Influence of Corporate Social Reputation and Brand Image on Customer Loyalty: The Mediating Role of Customer Satisfaction - A Study on Company A in the Galle District.

¹Nallaperuma, R.H.*, ²Ratnayake, G.S.
ruweena.20230949@iit.ac.lk, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

This study investigates the influence of corporate social reputation and brand image on customer loyalty, with customer satisfaction acting as a mediating variable. The research focuses on customers of Company A in the Galle District, Sri Lanka, within the premium wellness and skincare industry. A quantitative research approach was adopted, and data were collected from 264 respondents using a structured questionnaire. The data were analysed using SPSS through reliability, correlation, regression, and mediation analyses.

The findings reveal that both corporate social reputation and brand image have significant positive effects on customer satisfaction. Customer satisfaction emerged as the strongest predictor of customer loyalty, indicating its critical role in shaping repeat purchase intentions and positive behavioural outcomes. Corporate social reputation was found to have a significant direct effect on customer loyalty, demonstrating that consumers value ethical business practices and organisational credibility when developing long-term brand relationships. In contrast, brand image showed no significant direct effect on customer loyalty, suggesting that its influence is primarily indirect through customer satisfaction. Mediation analysis further confirmed that customer satisfaction significantly strengthens the relationships between corporate social reputation, brand image, and customer loyalty.

The study contributes to consumer behaviour literature by highlighting customer satisfaction as a key mechanism linking brand-related perceptions to loyalty outcomes. The findings imply that managers should prioritise customer satisfaction, strengthen socially responsible practices, and maintain a favourable brand image to enhance customer retention and long-term loyalty.

Keywords: Brand Image, Corporate Social Reputation, Customer Loyalty, Customer Satisfaction, Mediation

A Hybrid Ensemble Approach for Detecting AI-Generated Reviews in E-Commerce Platforms

¹Silva, A.*, ²Comester, R.

aeron.20221547@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The growing prevalence of AI-generated fake product reviews poses a significant threat to consumer trust in modern e-commerce ecosystems. With the increased accessibility of large language models (LLMs) such as GPT-4 and GPT-5, fraudulent sellers have begun using these tools at scale to generate reviews that are syntactically polished and difficult to distinguish from authentic human-written content. This erosion of review credibility hinders consumer decision-making, distorts market competition, and raises regulatory concerns for both platform operators and consumer protection agencies. This project presents the development, evaluation and deployment of a hybrid ensemble model capable of classifying product reviews as either AI-generated or human-written. A dataset of 35,579 multi-platform reviews was compiled using publicly available datasets and web-scraped data from four major e-commerce platforms - Amazon, Daraz, Temu and eBay - supplemented by the McAuley Amazon Review corpus, alongside an AI-generated review corpus produced under structured prompt engineering using multiple LLMs. Eight linguistically driven metadata features were engineered, including sentence burstiness, hedging ratio, modal verb density, adjective density, punctuation diversity, review length, Flesch Reading Ease, and the Automated Readability Index.

A three-component soft-voting ensemble - ExtraTrees, Random Forest, and DistilBERT - was developed, combining TF-IDF bigram vectorization and the metadata vector for the classical models, with end-to-end fine-tuning for DistilBERT. Class imbalance was addressed using SMOTE on the training split, and ensemble weights were assigned proportionally to each model's F1-score.

Over a held-out test set of 7,116 reviews, the ensemble achieved 98.86% accuracy, an F1-score of 0.9886, and a ROC-AUC of 0.9986, with an AI-class precision of 1.00, meaning no human-written reviews were misclassified as AI-generated. Individually, DistilBERT achieved 98.97% accuracy, followed by ExtraTrees (97.61%) and Random Forest (94.62%). However, an adversarial evaluation using ten "humanized" AI reviews containing slang and informal language reduced ensemble accuracy to 50%, highlighting a critical limitation of current detection approaches against deliberately disguised AI content.

The solution was deployed as a Gradio web application providing real-time review analysis with per-model confidence breakdowns and linguistic feature outputs. Key contributions of this project include a novel multi-platform, multilingual dataset incorporating Sinhala and Sinhala-English reviews, and a demonstration of the effectiveness of hybrid ensembles combining engineered linguistic features with deep learning for AI-generated review detection.

Keywords: AI-generated reviews; ensemble learning; e-commerce; multilingual NLP; transformer models

Investor 360: An Interactive Financial Analytics Platform based on Benjamin Graham's investment fundamentals for the Colombo Stock Exchange

¹Mendis, B.K.N*, ²Alwis, D.

nithara.20220640@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Abstract. Retail investors in emerging stock markets struggle with recurring problems rising from financial illiteracy, Scattered data access, and the lack of structured, locally standardized analytical tools. This paper shows Investor 360, a web based financial analytics platform developed for the Colombo Stock Exchange (CSE) in Sri Lanka. The system analyses a decade of historical financial data for 186 listed companies across 21 sectors, applying a dual scoring engine developed from Benjamin Graham's value investing principles. Two investor profiles are generated as a Defensive model that analyses ten years of financial data under strict stability standards, and an Enterprising model that applies a five years analytical data with flexible thresholds oriented towards growth. Both models generate a weighted score of 100 by evaluating indicators including earnings per share consistency, dividend payout, liquidity ratios, debt ratios, price-to earning, and margin of safety. The platform is implemented in Python using Streamlit, Pandas, and Plotly, with financial data stored in a JSON structure enabling efficient multi-year retrieval. A backward validation study applying the model to 2019 CSE data and comparing outcomes in opposition to performance recorded through 2025 confirms the system's dependability in identifying fundamentally strong long-term investments. Score distributions of 30 to 90 for the Defensive model and 40 to 95 for the Enterprising model shows clear discriminative power across the dataset. The resulting platform closes the gap between theoretical investment principles and accessible, interpretable decision support for CSE investors

Keywords: financial analytics, value investing, rule-based scoring, Benjamin Graham, investment decision making.

The Impact of Influencer Credibility on Purchase Intention among Millennial Consumers in the XXX Soap Brand through Brand Trust in Colombo District

¹De Silva, B.N.*, ²Ratnayake, G.S.

ninuri.20220014@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

This study focuses on the effect of influencer credibility on the purchase intentions of millennial consumers of XXX Soap Brand in the Colombo District, with brand trust examined as a mediating variable. Drawing upon Source Credibility Theory, the Theory of Planned Behavior (Ajzen, 1991), and Social Influence Theory (Salehuddin et al., 2023), the research explores the dimensions of influencer credibility including trustworthiness, expertise, and attractiveness (Liu and Zheng, 2024) and their influence on consumer perceptions and behavioral intentions. A quantitative research methodology was employed to test the proposed relationships (Ghanad, 2023). The findings indicate that influencer credibility exerts a significant and positive impact on purchase intention (Coutinho, Dias and Pereira, 2023; Shahid and Ikram, 2024). However, brand trust does not emerge as a significant predictor of purchase intention (Bogdan, Dospinescu and Dospinescu, 2025), nor does it mediate the relationship between influencer credibility and purchase intention (Herdiani et al., 2025). The results underscore the growing relevance of influencer-based communication in shaping consumer decision-making (Ampornklinkaew, 2025), particularly in the context of low-involvement personal care products (Garg and Bakshi, 2024). Ultimately, this study shifts the focus from traditional corporate brand building to highly dynamic, creator-centric consumer journeys (Sheryl et al., 2025). It proves that in contemporary digital environments, the immediate social proof of an individual influencer holds a far more agile and commanding influence over fast-moving consumer purchases (Migkos, Giannakopoulos and Sakas, 2025) than the institutional weight of the corporate brand itself. Consequently, these insights contribute to the literature on digital marketing and consumer behavior, while offering practical implications for marketers seeking to design effective influencer strategies in highly competitive digital markets.

Keywords: Brand Trust, Influencer Attractiveness, Influencer Credibility, Influencer Expertise

A Behavioral Clustering-Based Forecasting Framework for Sri Lankan Kernel-Based Coconut Exports

¹Elwitigala, V.A.*, ²Kirupananda, A.

anargi.20221262@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Sri Lanka ranks among the top five global coconut exporters. However, the Coconut Development Authority (CDA) lacks a data-driven forecasting system for export planning. This study presents a behavioral clustering-based forecasting framework for twelve kernel-based coconut export products from 2015-2025. Four data sources were merged into 1,572 monthly records. The two main Components in the study consists of: (1) Unsupervised clustering using K-Means, validated against Hierarchical, Gaussian Mixture, and Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithms with bootstrap stability (ARI=0.77), identified four distinct product behavior groups. (2) A range of forecasting models across baselines, classical time series AutoRegressive Integrated Moving Average (ARIMA), Seasonal AutoRegressive Integrated Moving Average (SARIMA) and Holt-Winters, Machine Learning (Random Forest, XGBoost), and Weighted Ensemble categories were evaluated per product.

Further to the forecasting it was confirmed that Machine learning outperformed time series for 10 of 12 products, with the highest accuracy of 89.43% received for Virgin Coconut Oil, using Random Forest. Furthermore, Validation against unseen 2025 actual data revealed systematic under-prediction in 8 of 12 products, attributed to structural market shifts confirmed by the Kolmogorov-Smirnov drift test. Furthermore, an Export Opportunity Score was brought into the research, combining forecast accuracy, growth trend, stability, and cluster profiles ranked Coconut Milk Powder as the top investment opportunity. Lastly, the system is deployed as a Gradio dashboard with MongoDB Atlas integration, which helps deliver actionable export intelligence for CDA.

Keywords: Coconut Exports, Time Series Forecasting, Machine Learning, K-Means Clustering, Export Opportunity Scoring, Sri Lanka

Track 2: Financial, Risk and Healthcare Analytics

A Personalised Data-Driven Approach to Predicting Brain Tumour Survivability in Adult Patients in Sri Lanka

¹Shazneen, H.*, ²Lelkada, P.

shazneen.20221504@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Brain tumours remain among the most aggressive and life-threatening forms of cancer, with survival outcomes varying considerably across patients. Accurate prognosis is essential for treatment planning and follow-up care; however, many existing predictive approaches lack interpretability, which reduces clinician trust and hinders their adoption in real-world healthcare settings. These challenges highlight the need for an explainable and clinically integrated framework for personalised survival prediction.

This study presents the development of an interpretable, web-based machine learning system for predicting individual survival outcomes in patients with brain tumours, with a focus on supporting clinicians in Sri Lanka. The objective was to provide personalised time-to-event survival predictions while delivering transparent and clinically meaningful explanations to support decision-making. The system integrates clinical, demographic, and histopathological variables within a comprehensive predictive framework.

A dataset containing records from over 1,400 patients with brain tumours was utilised. Multiple survival analysis techniques, including Cox Proportional Hazards, Aalen Additive Regression, Accelerated Failure Time (AFT), Random Survival Forest (RSF), and Explainable Boosting Machines, were evaluated using concordance index (C-index), Brier scores, and cross-validation. Unsupervised learning was also employed through K-Means clustering to identify patient archetypes and recovery patterns. Explainability was incorporated using SHAP values and global time-ratio interpretations derived from the AFT model.

The Random Survival Forest achieved the highest predictive performance with a test C-index of approximately 0.85. However, the AFT model was selected for deployment due to its strong predictive capability (test C-index $\approx$ 0.82-0.83), superior interpretability, stable calibration, and clinically meaningful outputs. Five patient archetypes were identified, and Kaplan-Meier survival curves supported effective risk stratification across patient groups.

The findings demonstrate the value of combining predictive accuracy with explainability in AI-assisted oncology. The proposed system offers a practical foundation for personalised prognosis generation, improved clinical decision support, and future development of patient-centred cancer care solutions.

Keywords: Brain Tumour, Clinical Decision Support System (CDSS), Explainable AI, Machine Learning, Survival Prediction

ForgeDetect: A Hybrid AI-Based Financial Document Forgery Detection System for Credit Lending Verification in Financial Institutions

¹Kaveesha, R.*, ²Abeykoon, J.

ruvini20212205@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Financial document fraud remains a significant challenge in credit lending operations, where customer-submitted documents such as bank statements and salary slips are frequently used to assess creditworthiness. According to the Entrust Identity Fraud Report, digital document forgeries account for approximately 57.46% of all document fraud cases, reflecting a 244% year-over-year increase and a 1,600% rise since 2021 (Entrust,2025). The growing accessibility of sophisticated image editing technologies has further increased the risk of document manipulation, creating challenges for financial institutions in accurately assessing creditworthiness and managing fraud-related risks. In Sri Lanka, verification of financial documents within lending processes is still dependent on manual review, making institutions vulnerable to forged submissions, operational inefficiencies, and inaccurate credit risk assessments.

This study proposes **ForgeDetect**, a hybrid AI-based financial document forgery detection system developed to support credit lending verification within financial institutions. A dataset comprising genuine and tampered financial documents was collected using the DocTamper benchmark dataset. The proposed methodology integrates Optical Character Recognition (OCR), copy-move forgery detection, localized Error Level Analysis (ELA), and a Two-Stream Convolutional Neural Network (CNN) to identify suspicious modifications. These forensic and learning-based signals are fused to generate explainable forgery assessments, confidence scores, and visual evidence highlighting suspicious regions. The solution was implemented as a full-stack web application using React, FastAPI, PyTorch, OpenCV, and EasyOCR. Furthermore, usability evaluation conducted with domain-relevant participants demonstrated positive acceptance of AI-assisted document screening, particularly in improving verification transparency and reducing manual effort. The findings indicate that a hybrid combination of deep learning, digital forensics, and document-aware validation techniques can provide an effective and compliance-aware decision-support mechanism for financial institutions while maintaining human oversight in credit lending decisions.

Keywords: Credit Lending, Digital Forensics, Financial Document Forgery Detection, Financial Institutions, Hybrid AI Systems

Reference: Entrust Cybersecurity Institute. (2025). 2025 identity fraud report. Available at: https://www.entrust.com/sites/default/files/documentation/executive_summaries/2025-fraud-identity-exec-summary.pdf

Debtor Tracking and Risk Management Using Machine Learning and Behavioural Analytics

¹Stephen, S.*, ²Liyanage, R.

shamendry.20222160@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Managing accounts receivable is an important part of maintaining stable cash flow within organisations. In many companies, debtor monitoring is still carried out using ageing reports and spreadsheet-based tracking methods. While these approaches help identify overdue balances, they often provide limited insight into future payment behaviour and make proactive risk management difficult. This project was developed to explore how machine learning and behavioural analytics can be used to support earlier identification of high-risk customers. The proposed system combines customer segmentation, predictive modelling, and explainable artificial intelligence techniques to create a more data-driven approach to debtor tracking and financial risk analysis.

The system was developed using anonymised accounts receivable and sales transaction data provided by a Sri Lankan tea manufacturing company operating in international markets. The dataset included 6,306 accounts receivable records across 656 customers, together with over one million sales transaction records spanning seven years. After extensive data cleaning and preprocessing, 24 behavioural and financial features were engineered using invoice activity, receipt behaviour, debit and credit note patterns, and historical purchasing trends.

Customer segmentation was performed using K-Means clustering to identify commercially meaningful behavioural groups, including Strategic High-Value customers and High Value One-Off buyers. These segment labels were then incorporated into a supervised classification model to improve predictive capability. Three machine learning algorithms were evaluated, namely Logistic Regression, Gradient Boosting, and Random Forest. Random Forest achieved the strongest overall performance, producing a cross-validated F1 Macro score of 0.8704 and a ROC-AUC score of 0.8645.

To improve transparency and support practical decision-making, SHAP explainability was integrated into the system to show the key behavioural and financial drivers behind each customer's risk score. The final outputs were deployed through an interactive Streamlit dashboard that enables finance teams to monitor portfolio exposure, analyse customer-level risk, and review segment-based financial patterns through visual analytics.

The results demonstrate that combining behavioural analytics with machine learning can provide a more proactive and explainable approach to debtor management compared to conventional ageing-based methods. The proposed system highlights the potential for predictive analytics to support financial decision-making and reduce exposure to high-risk accounts before payment defaults occur.

Keywords: behavioural analytics, credit risk, explainable AI, machine learning, receivables management

SVEMS: AI-Assisted Software Vendor Evaluation Management System for Transparent and Data-Driven Software Selection

¹Gunawardana, D. I. A.*, ²Zeezan, A.

dinangaguna@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka.

Selecting the wrong software vendor is not only a procurement mistake; it can trigger implementation failure, operational disruption, financial loss, and long-term strategic setbacks. This risk is especially critical for small and medium-sized enterprises, where vendor selection is often carried out through fragmented manual practices such as spreadsheets, email threads, informal discussions, and subjective judgement. These approaches weaken transparency, reduce traceability, limit evaluator collaboration, and make it difficult to justify final decisions with consistent evidence. Although existing procurement and vendor management platforms support activities such as sourcing, contract management, and post-selection monitoring, scant attention is given to the pre-selection stage, where vendor suitability must be evaluated before implementation risk becomes business impact. This study presents SVEMS, an AI-assisted Software Vendor Evaluation Management System developed to support transparent, structured, and evidence-based vendor selection while preserving human evaluator control. The system centralizes the pre-selection lifecycle through project-based evaluation, requirement baseline management, vendor submissions, evaluator collaboration, weighted scoring, milestone tracking, dashboards, reports, and role-based governance. Its AI-assisted capabilities support requirement-proposal matching, cost extraction, CAPEX/OPEX and total cost of ownership analysis, vendor ranking, anomaly and bias detection, evaluation summaries, and chatbot-based guidance. Rather than replacing evaluator judgement, SVEMS follows a human-in-the-loop approach where AI outputs are reviewed alongside manual scores, evaluator comments, and approvals. This improves analytical depth while maintaining accountability and professional judgement.



Figure 1: SVEMS AI-Assisted Vendor Evaluation Workflow

The developed prototype demonstrates how AI-assisted decision support can transform software vendor evaluation from a scattered and judgement-driven process into a structured and evidence-based selection. By centralizing evaluation records, improving workflow visibility, enabling requirement-level proposal comparison, and strengthening accountability, SVEMS supports strongly defensible vendor decisions. Initial evaluation suggests that the system can reduce reliance on subjective judgement, improve collaboration among stakeholders, and enhance confidence before implementation. This research contributes to a practical SME-focused decision-support framework and highlights the value of artificial intelligence in enabling fairer and more reliable software vendor evaluation outcomes.

Keywords: Artificial Intelligence, Decision Support System, Requirement-Proposal Matching, Software Vendor Evaluation, Total Cost of Ownership

CeyFleet: Enhancing operational and financial efficiency through effective crew management and skipper recommendation in Sri Lanka's multi-day fishing vessels.

¹Welikala, D.T.*, ²Abeykoon, J.

welikaladevmi@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The fisheries industry is a vital pillar of the Sri Lankan economy. However, inefficiencies in vessel and crew management on commercial fishing vessels undermine the success of the harvest, revenue, and the industry's overall performance. These business owners rely on personal memory and judgment when making critical decisions due to limited access to data and insights. This project aims to develop “CeyFleet”, a solution tailored to enhance operational and financial efficiency through effective crew management and skipper recommendations for commercial multi-day fishing vessels operating in Sri Lanka.

A hybrid agile project management approach was adopted throughout the project to support incremental development. Interviews were conducted to gather requirements, and several stakeholders were followed up with to incorporate their expertise and feedback into the application design. The system architecture is a Progressive Web Application built with Next.js, using PostgreSQL as the database. OTP-based authentication is used to access the application. Skipper performance records are stored in a centralized database that generates AI-driven recommendations using OpenAI, and Cloudinary was utilized for document storage.

CeyFleet replaces subjective decision-making with structured data and actionable insights, thereby enabling smarter skipper recruitment, improved crew management, and operational transparency. The platform enhances vessel efficiency and reflects a broader shift towards digital transformation in Sri Lanka's commercial marine fisheries industry.

Keywords: Commercial multi-day fishing vessels, AI-generated skipper recommendation, Vessel management, Document management, Data-driven insights

Automated Coco Peat Block Quality and Defect Detection System for Improved Production Efficiency


¹Muniprema, R.*, ²Ratnayake, G.

ramudi.20210846@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Sri Lanka is one of the world's leading exporters of coco peat products, which are widely used in horticulture and sustainable agriculture due to their excellent water retention and growing properties. Maintaining consistent product quality is essential to meet international export standards; however, quality inspection within the coco peat manufacturing industry is still largely dependent on manual visual assessment. This process is time-consuming, subjective, and prone to human error, resulting in inconsistent quality evaluation, increased production inefficiencies, and potential export rejections.

This research proposes an IoT-enabled quality inspection system for automated coco peat block quality assessment and defect detection. The primary objective is to improve the accuracy, consistency, and efficiency of quality control operations while reducing dependence on manual inspection methods. The proposed system integrates an ESP32-CAM module, TCRT5000 infrared sensors, an L298N motor driver, and a conveyor-based inspection mechanism to perform automated real-time quality assessment. Computer vision and image processing techniques are employed to capture and analyse images of coco peat blocks during production. OpenAI-based image analysis models are utilized to identify visual defects and quality inconsistencies, while Firebase provides cloud-based data storage and monitoring capabilities. The prototype was developed using Python, Embedded C/C++, HTML and CSS with communication facilitated through Wi-Fi and HTTP-based IoT technologies.



A functional prototype has been successfully developed to demonstrate the feasibility of integrating machine vision, machine learning, and IoT technologies within the coco peat manufacturing process. Preliminary observations indicate that the proposed system can improve inspection consistency, support real-time decision-making, and reduce operator dependency when compared to conventional manual inspection methods. The proposed solution contributes to Industry 4.0 adoption within Sri Lanka's coco peat

industry by providing a scalable, cost-effective, and intelligent quality assurance system capable of enhancing productivity, maintaining international quality standards, and strengthening global competitiveness.

Keywords: Coco Peat, Computer Vision, Defect Detection, IoT, Manual Inspection

Intelligent Chronic Care Platform - Medikeep

¹Dodawatte D.*, ²Balathasan Y.

dulmini.20212103@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Healthcare data management remains a significant challenge in Sri Lanka, particularly for patients with chronic diseases who require continuous monitoring and long-term care. Medical records are often scattered across paper documents, emails, messaging platforms, and multiple healthcare providers, resulting in fragmented information, delayed decision-making, repeated diagnostic tests, and inefficient healthcare delivery. These issues affect both patients and clinics by reducing continuity of care, increasing operational costs, and limiting access to complete medical histories during critical situations.

To address these challenges, this project proposes Medikeep, an intelligent chronic care management platform designed for patients and clinics. The system enables secure digital management of medical records, long-term tracking of patient history, and controlled sharing of healthcare information. Patients can store records from multiple providers, share selected documents securely, access AI-generated explanations to better understand medical reports, and use an AI-powered chatbot to quickly retrieve specific information from their medical records. Clinics can efficiently manage patient data, access comprehensive medical histories, and retrieve critical information during emergencies.

The platform was developed using Next.js, FastAPI, PostgreSQL, and OpenAI-powered AI services, incorporating OCR-based document processing and a Retrieval-Augmented Generation (RAG) chatbot for intelligent information retrieval. The chatbot retrieves relevant information from patient-specific medical records, grounds responses using retrieved document content, provides supporting references, and incorporates confidence-based retrieval and hallucination reduction techniques to improve response reliability. Security is maintained through JWT authentication, role-based access control, audit logging, and secure sharing mechanisms.

Requirements were identified through literature review, patient surveys, clinic interviews, and analysis of real-world healthcare scenarios. The system was evaluated through functional testing and user feedback from patients. Results indicate that Medikeep improves record accessibility, continuity of care, patient understanding of medical information, and workflow efficiency. The RAG-based chatbot demonstrated improved relevance and reliability compared to generic AI responses by grounding answers in patient-specific medical records. Overall, the system demonstrates the potential to support more connected, secure, and data-driven healthcare delivery while remaining practical for adoption in small and medium-sized healthcare settings.

Keywords: Artificial Intelligence, Chronic Disease Management, Electronic Health Records, Healthcare Data Sharing, Retrieval-Augmented Generation

Automated Personal Loan Underwriting Using Integrated Rule Based and Behavioral Analytics

¹Seneviratne, S.K.*,²Samarasekera, N.

kavindya.20221746@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

An automated personal loan underwriting system was developed by integrating rule-based validation, machine learning, customer segmentation, and loan decision making within a single end-to-end platform. Using data from a leading commercial bank in Sri Lanka, the system addresses key challenges in traditional loan underwriting, including lengthy approval times, inconsistent assessments, and underutilization of internal customer behavioral data. While developed using a specific institutional setting, the proposed methodology is transferable to comparable retail banking environments.

A fully anonymized dataset of 2,509 customers was extracted from internal banking systems in compliance with Sri Lanka's Personal Data Protection Act. Customer details, account records, transaction histories, loan cashflows, and repayment data were integrated to construct a comprehensive analytical framework. The technical complexity of the project lies in its integration of multiple interconnected components rather than reliance on a standalone predictive model.

The system first applies a rule-based eligibility engine that validates applicants against regulatory and operational requirements, including age restrictions, active banking relationships, employment verification, and financial capacity assessment. Customers with insufficient observable financial history were identified through a thin-file process and routed for manual review.

For credit default prediction, three machine learning approaches such as CatBoost, XGBoost, and XGBoost with monotone constraints were systematically evaluated. Each model was tested using both a baseline and a hyperparameter tuned version, producing six model variations for comparison. Feature engineering was performed using behavioral, demographic, transactional, and repayment related variables while preventing target leakage. The final model was evaluated using the Receiver Operating Characteristic Area Under the Curve (ROC-AUC), precision, recall, F1-score, and confusion matrix analysis.

Customer segmentation was further conducted using the K-Prototypes clustering algorithm to group customers into behavioral profiles that improve interpretability in lending decisions. Finally, the system was deployed as two interconnected web applications consisting of a customer facing eligibility portal and an internal officer dashboard, demonstrating a practical and deployable solution for faster, more transparent, and consistent loan underwriting. Future work may focus on integrating external credit bureau data, strengthening role-based access control mechanisms, and a pathway to onboard customers who do not yet hold accounts with the bank.

Keywords: Automated Loan Underwriting, Credit Risk Modelling, Feature Engineering, Machine Learning, Rule-Based Systems

Brain Tumour MRI Classification Using Ensemble Deep Learning and Explainable AI

¹Seneviratne, T.M*, ²Hassan, F.

maliksha.20231328@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Brain tumour classification from MRI images is a clinically significant task, yet deep learning models used for this purpose are often criticised for their lack of transparency. This study presents an explainable ensemble deep learning framework for four-class brain tumour MRI classification, distinguishing glioma, meningioma, pituitary tumour, and normal brain images.

Two pre-trained convolutional neural network architectures, ResNet50 and EfficientNet-B0, were adapted using transfer learning on the PMRAM Bangladeshi Brain Cancer MRI Dataset available through Mendeley Data, and subsequently fine-tuned through Optuna-based hyperparameter optimisation. Base model accuracies of 86.27% and 84.97% for ResNet50 and EfficientNet-B0 respectively were improved through fine-tuning, with ResNet50 reaching 96.73% after task-specific adaptation. A key observation during this stage was the recurring confusion between pituitary and normal classes in single-slice classification, which was substantially reduced following fine-tuning. To further improve performance, multiple ensemble strategies were investigated. Probability-based methods, including weighted soft voting, confidence thresholding, and entropy thresholding, performed poorly and in some cases reduced accuracy, whereas class-wise weighted logit averaging with grid-search-optimised weights produced the strongest results. The final tuned class-wise weighted logit ensemble achieved a test accuracy of 99.35%, demonstrating that combining fine-tuned models through logit-level fusion significantly outperforms both individual models and probability-based combination approaches.

To address the interpretability of model predictions, three post hoc explainability techniques were applied and compared: Grad-CAM, Grad-CAM++, and Score-CAM. Grad-CAM was found to provide the clearest and most anatomically meaningful heatmaps, highlighting tumour-relevant regions in the MRI images, and was selected for integration into a Streamlit-based interactive dashboard. The dashboard allows users to upload MRI scans, receive class predictions with confidence scores, and inspect Grad-CAM visualisations, combining high classification performance with practical interpretability in a single deployable prototype.

Keywords: Brain Tumour Classification, Deep Learning, Ensemble Learning, Explainable AI, MRI Imaging.

A Data-Driven Approach to Predicting Hospital Length of Stay, Mortality and Readmission Risk in Myocardial Infarction Patients

¹Awn, M.*, ²Hassan, F.

awn.20222469@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatic Institute of Technology, Sri Lanka

Myocardial infarction (MI) remains a leading cause of hospitalisation, yet many machine-learning risk models rely on detailed laboratory and vital-sign data that are not consistently available in structured form in many healthcare settings. We ask whether routinely collected administrative hospital data alone can support clinically meaningful, interpretable risk stratification. Using 16,537 MI admissions (13,152 patients) from the MIMIC-IV database, we develop a two-stage framework to predict four outcomes: length of stay (LOS), admission mortality, post-admission mortality and 30-day readmission risk. A strict patient-level train-test split was applied to prevent data leakage, and model performance was assessed using ROC-AUC with confidence intervals. The best-performing models achieved a ROC-AUC of 0.829 (95% CI: 0.815-0.843) for LOS and 0.820 (0.798-0.840) for post-admission mortality. Including diagnosis-related group (DRG) variables improved mortality prediction, increasing ROC-AUC from 0.747 at admission to 0.820 post-admission. SHAP analysis showed that comorbidity burden was consistently one of the most important predictors across all models. DRG severity and mortality scores were particularly important for post-admission mortality, reflecting changes in patient condition during the hospital stay. A Cox proportional hazards model further identified prior MI, emergency admission and discharge to home health care as key factors associated with earlier readmission. The models were integrated into a low-cost clinical decision support dashboard and tested using contrasting patient profiles. These findings show that administrative hospital data contains useful predictive signal across multiple MI outcomes and can support practical, data-driven decision-making, especially in settings where detailed clinical data are limited.

Keywords: myocardial infarction, clinical decision support, electronic health records, interpretable machine learning, class imbalance

Track 3: Language Models and Explainable AI

Mechanistic Interpretability Under Model Compression for Transformer-Based Language Models

¹Faumi, N.*, ²Sumanathilaka, D.

nabeelahfaumi@gmail.com, ¹Robert Gordon University, United Kingdom

²Informatics Institute of Technology, Sri Lanka

Transformer-based models are increasingly compressed to reduce computational costs, yet the impact on their internal architecture remains poorly understood. This study investigates how weight pruning, post-training quantisation (PTQ), and knowledge distillation affect the mechanistic interpretability of GPT-2 Small.

By evaluating the model against known ground truth circuits Indirect Object Identification (IOI) and the Greater-Than task, this research identifies three distinct mechanistic failure modes. Magnitude attention head pruning produces zeroing, the physical elimination of circuit components while accuracy is maintained through residual non-circuit computation. Magnitude neuron pruning produces inversion, in which intact attention heads are redirected toward incorrect outputs due to a corrupted residual stream. INT8 quantisation produces dampening, a sign-preserving attenuation of circuit contributions that preserves task accuracy while reorganising causal pathways.

Findings indicate that coarse metrics such as cosine similarity and logit lens analysis systematically underestimate internal damage. Sparse autoencoder (SAE) analysis revealed catastrophic early-layer feature disruption invisible to standard diagnostics. The INT8 Greater-Than paradox further demonstrates that near-baseline task accuracy can coexist with circuit-level reorganisation comparable to 70% pruning, establishing that output-level evaluation is insufficient for assessing circuit integrity. Knowledge distillation revealed a structural asymmetry: MLP-dominant circuits survived depth reduction through functional replication, while attention-dominant circuits collapsed. These results establish that compression qualitatively reshapes internal logic in method-specific ways, necessitating circuit-level validation for the safe deployment of compressed language models

Keywords: Circuit Analysis, Features, Mechanistic interpretability, Model compression, Pruning

SINBRIEF: An Unsupervised Framework for Abstractive Text Summarisation of Sinhala Legal Documents

¹Lasandi, M.*, ²Jayatilleke, N.

minduli.20220374@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka,

Abstractive summarisation aims to generate concise and coherent summaries that capture the essential information of a document using newly formed sentences rather than directly extracting text. In Sri Lanka, legal documents such as Acts and Bills are often lengthy, complex, and written in domain-specific language, making them difficult to interpret for both legal professionals and the general public. Despite the growing importance of legal information accessibility, research in Sinhala Natural Language Processing (NLP), particularly within the legal domain, remains limited due to the scarcity of linguistic resources and annotated datasets. This research presents SinBrief, an unsupervised abstractive summarisation framework for Sinhala legal documents that operates without the need for large-scale labelled data. The proposed approach combines sentence clustering to identify thematic structures, domain-aware weighted word graph construction to model relationships between legal concepts, and graph-based sentence fusion to generate candidate abstractive summaries. To assess the quality of generated candidates, a sentence scoring model is trained and employed to rank sentences based on their relevance and coherence, enabling the selection of informative summary content while preserving document context. To support research in this low-resource domain, this study also introduces SinhaLegal, a comprehensive dataset comprising Sinhala legal documents collected from legislative sources. Furthermore, a domain-adapted language model based on Llama 3.1 is developed through continual pre-training on Sinhala legal text, enhancing the representation of legal terminology and contextual knowledge. Evaluation using reference-free summarisation metrics indicates that the proposed framework generates more abstractive and semantically consistent summaries compared to both extractive and abstractive baselines, while maintaining competitive performance across factual consistency metrics. By integrating graph-based abstraction, sentence scoring, and domain adaptation, the proposed framework produces structured summaries that effectively capture information across legal documents while reducing reliance on extensive annotated resources. This work contributes to the advancement of Sinhala legal NLP by introducing an unsupervised abstractive summarisation framework, releasing valuable language resources, and facilitating improved access to complex legal information for researchers, legal practitioners, and the wider public.

Keywords: Abstractive Summarisation, Domain Adaptation, Low-Resource Languages, Sinhala Legal Documents, Unsupervised Learning

CAL-Log: Cost-Aware Active Learning with Logarithmic Cognitive Effort Modeling and Online Adaptation to Human Annotation Behavior

¹Kariyakaranage, V.S.*, ²Athuraliya, B.

vihangasupasan2001@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Active Learning (AL) reduces the amount of labeled data required to train Natural Language Processing (NLP) models, yet most existing methods optimize label efficiency while ignoring the human cost of annotation. Standard uncertainty sampling implicitly assumes that annotating a 500-word document requires the same effort as a 50-word tweet, and pipelines that attempt linear cost normalization suffer from cost collapse, a failure mode formally identified in this work, where selection becomes systematically biased toward short, low-information samples. Combined with poorly calibrated model confidence and static cost assumptions, this leads to inefficient use of limited annotation budgets in real-world labeling campaigns.

This research introduces CAL-Log, a cost-aware AL framework grounded in psycholinguistic reading theory. The acquisition function normalizes Shannon entropy by a logarithmic cost model, $C(x) = \alpha + \beta \log(1 + L(x))$, capturing the sub-linear relationship between document length and human reading effort. Expected Calibration Error is monitored each round, triggering a diversity-based fallback using SBERT embeddings when model confidence becomes unreliable, and an online Ordinary Least Squares regression continuously adapts the cost parameters to each annotator's observed behavior. A full-stack interactive simulation platform with a transparency dashboard was also developed, allowing researchers to observe live accuracy curves during annotation sessions.

CAL-Log was evaluated against seven baselines, Random, Entropy, Least Confidence, Margin, CoreSet, BADGE, and CAL-Linear, across ten benchmark text classification datasets under a 120-minute annotation budget. CAL-Log achieved an AUC-F1-Cost of 0.6096, outperforming Margin by 29% and BADGE by 30.6%, and reached $F1 = 0.80$ in 38.3 minutes compared to 93.7 minutes for Random and 148.5 minutes for Entropy, an up to $3.88\times$ speedup, with a $5.52\times$ speedup on long-document datasets such as Yelp Polarity. Statistical significance was confirmed ($p = 0.031$) with a large effect size (Cohen's $d = -1.21$). A live deployment with preliminary user evaluation ($N = 42$) further suggests that the online cost model produces reading-speed classifications consistent with annotator self-reports. These findings demonstrate that modeling annotation cost realistically, rather than assuming uniform effort, substantially reduces wasted annotator workload and offers a practical, generalizable path toward more efficient human-in-the-loop NLP systems.

Keywords: Active Learning, Cost-Aware Annotation, Human-Centered NLP, Regret-Safe Learning, Text Classification

Apofasi: A Unified Temporally Aware Retrieval-Augmented Generation Framework for Question Answering across Statutory and Case Law in Evolving Mixed Legal Systems

¹Sathurshanan, M.*, ²Athuraliya, B.

shathurshmanoharan@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Current retrieval augmented generation (RAG) frameworks, despite being powerful for general tasks, face significant challenges when applied to law. This is fundamentally due to their temporal blindness, which makes them incapable of handling the complexities of an evolving legal system. Standard RAG architectures are unable to differentiate between different historical versions of a statute or precedent, leading to factually invalid answers. This problem is even more severe in mixed legal systems, such as Sri Lanka's, where statutory and case law are interconnected. Such jurisdictions require a more comprehensive reasoning framework than what current systems can provide.

To overcome these constraints, this research proposes the Apofasi framework, a temporally aware RAG architecture. The important part of this methodology is the design and construction of a unified temporal knowledge graph (TKG), which models the evolution of both statutory law and case law. Querying is managed by a strategic, time conditioned hybrid retrieval mechanism that combines neural vector search for semantic relevance, symbolic graph traversal for structural legal authority, and a deterministic temporal drift analyser to enforce strict chronological validity, ultimately generating legally grounded and reliable answers.

The proposed framework, Apofasi, demonstrates substantial improvements over standard semantic and graph-based RAG baselines. Most notably, the system achieves an outstanding Temporal Precision@5 score of 0.8560, significantly outperforming the standard semantic RAG (0.3094) and pure graph-based RAG (0.6221) baselines. These findings validate the framework's ability to effectively mitigate temporal hallucinations, ensuring that the generated answers are not only semantically relevant but also strictly legally valid for the queried time.

Keywords: Legal Question Answering, Natural Language Processing (NLP), Retrieval-Augmented Generation (RAG), Temporal Knowledge Graph

XFedRec: A Robust, Adaptive, and Explainable Federated Recommendation Framework

¹Ravichandran, K.*, ²Poravi, G.

keshav.20200326@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Federated Learning (FL) provides a decentralized, privacy-preserving paradigm for modern recommendation systems. However, real-world e-commerce environments expose FL to compounded vulnerabilities. A primary challenge is the Security-Agility Trade-off: robust aggregation methods designed to neutralize Byzantine attacks mathematically suppress statistical outliers, inadvertently preventing the system from adapting to genuine Concept Drift. Furthermore, the inherent opacity of deep neural networks complicates compliance with transparency mandates such as GDPR Article 22.

In this paper, we propose XFedRec, a unified framework that synthesizes robust aggregation, concept drift adaptation, and edge-based Explainable AI (XAI). Utilizing a Neural Collaborative Filtering (NCF) backbone, XFedRec introduces a client-side Hybrid Concept Drift Detector combining a 1D Kalman Filter, Z-Score sudden triggers, and ADWIN2 to isolate and adapt to shifting preferences amidst highly skewed Non-IID data. Server-side, Coordinate Median aggregation, complemented by Top-K Sparsification (80 percent compression) and Differential Privacy (epsilon of 2.0), neutralizes 10x Byzantine scaling attacks.

Experiments on a Dirichlet-partitioned (alpha of 0.1) MovieLens-100K dataset demonstrate that XFedRec achieves near-zero drift detection delay and rigorously recovers continuous rating accuracy (RMSE). Crucially, our ablation studies empirically formalize the Security-Agility trade-off, demonstrating how rigid robust aggregators inherently dampen adaptation velocity in dynamic federated networks.

Keywords: Byzantine Robustness, Concept Drift, Differential Privacy, Explainable AI, Federated Learning

Autodistil-KG: Automated LLM Distillation with Knowledge Graph Supervision

¹Periyapperuma, N.*, ²Wijesinghe, N.

nimesh.20221000@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Problem: Large Language Models (LLMs) demonstrate remarkable generative capabilities but remain prone to hallucination in domain-specific applications. While Graph Retrieval-Augmented Generation (GraphRAG) mitigates this by grounding generation in structured knowledge, it introduces significant inference latency due to mandatory runtime graph traversal and couples the deployed system to a potentially large graph database, thereby imposing significant operational overhead. **Methodology:** This paper presents *AutoDistil-KG*, an automated, end-to-end framework that shifts the computational burden of knowledge-graph grounding from inference to training through a four-stage pipeline: (1) *Graph Traverser* - performs multi-hop traversal of the knowledge graph using an LLM-guided Reasoning strategy, where the teacher model actively selects traversal paths based on relational context, synthesising rich subgraph analyses, (2) *ChatML Formatter* - converts synthesised question-answer pairs into a supervised fine-tuning dataset in ChatML conversational format, (3) *Finetuner* - applies 4-bit QLoRA-based Parameter-Efficient Fine-Tuning (PEFT) via Unsloth to transfer the distilled knowledge into a smaller student model and (4) *Evaluator* - benchmarks the distilled student against the base model and a GraphRAG baseline under identical conditions using LLM-as-a-judge (G-Eval) and n-gram metrics.

Evaluation: Evaluated on HETIONET, the distilled Qwen3.5 4B student model achieves a judge-average score of 0.524, representing a 32.0% improvement over the base model (0.397) and a 164.6% improvement over GraphRAG (0.198). Scope grounding - measuring whether the model stays within the knowledge graph's boundaries - improves from 0.620 (base) to 0.760 (distilled), compared to 0.233 for GraphRAG. ROUGE-L improves by 39.3% and BLEU-4 by 152.3% over the base model, confirming lexical fidelity to the teacher-generated reference answers. At the system level, the distilled model achieves an average inference latency of 10.1 s, approximately 15× faster than GraphRAG (151.8 s) and 32% faster than the base model (14.8 s), while producing 32% fewer output tokens, indicating more concise and accurate responses without any inference-time computational penalty.

Contributions: The principal contributions of this work are: (i) a fully automated pipeline integrating graph traversal, dataset synthesis, LoRA fine-tuning, and multi-system evaluation into a single configurable framework, (ii) empirical evidence that BFS-based multi-hop traversal produces training data sufficient for a small student model to outperform both its base counterpart and a GraphRAG baseline and (iii) a reusable multi-system evaluation methodology for comparing distillation-based and retrieval-based approaches under controlled conditions. These results provide directional evidence that pre-encoding structured knowledge graph context into model parameters during training is a viable alternative to inference-time retrieval for domain-specific question answering.

Keywords: Graph Retrieval-Augmented Generation, Knowledge Distillation, Knowledge Graphs, Large Language Models, Parameter-Efficient Fine-Tuning

Quantarra: An XAI Augmented Automated Recommendation Framework for Multi Issue Data Quality Detection in Machine Learning Pipelines

¹Gamage, T.*, ²Caucidheesan, K.

w1902266@my.westminster.ac.uk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The predictive validity of machine learning models is fundamentally constrained by the quality of their training data, yet practitioners lack a unified, domain agnostic means of diagnosing the defects that most degrade it. Although each quality issue is well studied in isolation, no existing framework concurrently addresses all five canonical dimensions (missing values, outliers, duplicate records, label noise, and class imbalance), couples' detection with transparent instance level justification, and translates its findings into directly executable remediation. This study introduces Quantarra, a domain agnostic and explainability augmented framework that unifies these capabilities in a single automated pipeline. Outlier detection cascades Isolation Forest scoring, statistical pruning, and Local Outlier Factor refinement; duplicate detection applies a support vector classifier over prefix blocked candidate pairs; and label noise is identified through a Cross Validated Committees Filter. The heterogeneous outputs are normalized into a single interpretable Data Quality Score, while TreeSHAP, SHAP, and LIME explain each flagged record. A GPT-4o-mini engine, conditioned on privacy preserving issue metadata, then synthesizes dataset specific, executable remediation code, closing the loop between diagnosis and intervention.




Figure 1: Quantarra three tier system architecture: detection modules, XAI layer, recommendation module, and LLM

The framework was evaluated on five public benchmark datasets from the healthcare and finance domains, together exhibiting all five quality dimensions across distinct problem settings. The detection modules attained average F1 scores of 0.8325 for outliers, 0.7402 for duplicates, and 0.8141 for label noise, each surpassing the acceptability threshold and generalizing across domains rather than overfitting. Controlled ablation confirmed that the pruning stage and neighborhood calibration dominate outlier detection accuracy, and that the precision recall trade off in duplicate detection is governed by the classification threshold. Comparative analysis established that no existing tool simultaneously delivers full

automation, five-dimension coverage, integrated explainability, and language model generated remediation, positioning Quantarra as the first framework to unify them and a reproducible basis for trustworthy, explanation driven data preparation.

Keywords: Anomaly Detection, Data Quality, Explainable AI, Large Language Models, Machine Learning

HEXEval: A Holistic Framework for Standardizing Technical and Human-Centred Explainable AI Evaluation

¹Saheerdeen, S. F.*, ²Jawwadh S.

salman.20221380@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Explainable Artificial Intelligence (XAI) has become essential for increasing transparency and trust in machine learning systems deployed in high-stakes domains such as healthcare and finance. Although numerous explanation methods, including SHAP, LIME, Anchors and DiCE, have been proposed, there is no widely adopted and standardized approach for evaluating the quality of their explanations. Existing evaluation approaches are fragmented across technical and human-centred paradigms, employ inconsistent metrics, and often lack practical mechanisms for reproducible assessment. Consequently, selecting the most appropriate explanation method for a given stakeholder or application remains a significant challenge.

This research proposes HEXEval, a holistic evaluation framework that standardizes post-hoc XAI evaluation by integrating technical performance metrics with human-centred assessment criteria within a unified evaluation pipeline. The framework evaluates explanations using quantitative dimensions such as fidelity, parsimony, applicability and stability, alongside qualitative dimensions including interpretability, completeness, trust, usefulness, satisfaction and decision support. HEXEval further introduces an LLM-based persona evaluation mechanism that simulates stakeholder perspectives and enables scalable human-centred assessment. A recommendation engine combines evidence from both evaluation streams to generate stakeholder-specific rankings of explanation methods.

The framework was implemented and evaluated using healthcare and credit-risk prediction case studies involving SHAP, LIME, Anchors and DiCE explanations. Experimental results demonstrated that explanation methods exhibit different strengths and weaknesses across technical and human-centred dimensions, with no single method consistently outperforming others. Furthermore, validation against responses from 35 human participants across six Likert-scale dimensions showed that the LLM-based persona proxy produced ratings that closely aligned with human assessments, supporting its potential as a scalable evaluation alternative. These findings demonstrate the value of combining technical and human-centred evidence when assessing explanation quality.

HEXEval contributes a practical and reproducible framework that bridges the long-standing divide between technical benchmarking and human-centred XAI evaluation, enabling more informed and stakeholder-aware selection of explanation methods.

Keywords: Explainable Artificial Intelligence, Human-Centred Evaluation, Machine Learning, Trustworthy AI, XAI Evaluation.

An Interpretable Hybrid Ensemble Framework for Consumer Goods Supply Chain Analytics

¹Silva, V.S.*, ²Rathnayaka, R.M.P.S.

vinukassilva@gmail.com, ¹Robert Gordon University, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Consumer goods supply chains operate under volatile demand, compressed lead times, and rising cost pressure, yet most analytics deployments remain functionally fragmented, opaque to planners, and misaligned with asymmetric business costs. Demand forecasting, inventory optimization, return risk prediction, and anomaly detection are typically developed as isolated systems, so forecast uncertainty is never propagated to safety stock decisions, return risk is decoupled from net demand, and high-performing black-box models fail to earn practitioner trust. This research addresses these barriers through a modular Hybrid Artificial Intelligence (AI) framework that unifies all four analytics functions within a single, explainable, production-ready architecture. Following the Design Science Research methodology, the framework was implemented and empirically evaluated on three industry-standard benchmarks. Demand forecasting employs a multi-paradigm ensemble combining Seasonal AutoRegressive Integrated Moving Average (SARIMA), eXtreme Gradient Boosting (XGBoost), Long Short-Term Memory (LSTM), and attention-based Sequence-to-Sequence models under a Ridge Stacking meta-learner trained on out-of-fold predictions. Downstream modules include an asymmetric Newsvendor inventory optimizer with adaptive uncertainty estimation, a stacking ensemble for return risk classification, and seven comparatively evaluated anomaly detectors. Per-prediction SHapley Additive exPlanations (SHAP) are delivered alongside every forecast through a FastAPI REST interface at sub-second latency. Empirical results validate the framework across its entire pipeline. At its core, the Ridge Stacking meta-learner achieved a weighted Mean Absolute Percentage Error of 4.45% - a 66.5% improvement over the ARIMA baseline and a statistically significant gain over the strongest individual model, validated across 56 pairwise significance tests. Notably, the meta-learner assigned a negative coefficient to LSTM, revealing that structurally weaker models can contribute corrective information through unconstrained meta-learning. These forecasts, propagated downstream, drove inventory policies to a 99.6% service level with bullwhip ratios below 1.0 - while exposing a practically significant decoupling of forecast accuracy from inventory cost efficiency. Complementing the forecasting core, return risk classification reached a ROC-AUC of 0.7169, feeding net demand estimation, and Seasonal-Trend decomposition residual anomaly detection safeguarded data integrity with a point-adjusted F1 of 0.894, outperforming deep learning alternatives. Ablation studies further confirmed that selective Google Trends integration outperforms bulk external signal augmentation. All capabilities are further consolidated into a single interactive dashboard, giving planners explainable, real-time insight at a glance. These results confirm that statistical rigor, interpretability, and operational deployability can coexist within one unified supply chain analytics system, bridging the long-standing gap between predictive accuracy and trustworthy, cost-aligned decision support for consumer goods practitioners.

Keywords: Demand Forecasting, Ensemble Learning, Explainable AI, Inventory Optimization, Supply Chain Analytics

FGNN: Joint Flow-based Model-level Explainer for Heterogeneous GNNs with Node-Features

¹Minosha, J.*, ²Weeransinghe, R.

minoshanj@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Heterogenous Graph Neural Networks (HGNNs) underpin high-stakes applications such as recruitment platforms, academic collaboration networks, and biomedical knowledge graphs, yet they remain black-box models whose decisions cannot be explained or audited. EU AI Act and GDPR now mandate automated decision systems to be interpretable, shifting explainability from a desirable feature to a regulatory requirement. Model-level explanation methods address this by generating a synthetic graph that maximizes a GNN's predicted score for a target class. However, existing model-level explainers focus predominantly on homogeneous graphs. The few that address heterogenous graphs generate graph structure and node-features separately which yields semantically inconsistent explanations and require a separate trained generator for each graph size and node type.

This work introduces FGNN, the first model-level explainer to apply flow matching to heterogenous GNNs. A single class-conditioned heterogenous graph transformer jointly generates graph structure, node types, and node features in one forward pass, capturing the natural interdependencies that decoupled generators overlook. A single model generates graphs of any size or class, and schema validity is enforced within the architecture, thereby eliminating the need for both separate generators and post-processing correction. The model is trained per dataset on two-hop subgraphs and evaluated on the DBLP citation and IMDB movie networks.

FGNN matches or exceeds the state-of-the-art DiGNNExplainer on predictive faithfulness (0.999 on DBLP, 0.928 on IMDB) and ground-truth faithfulness (0.933 on IMDB), while using 300 times fewer model forward passes per candidate and achieving an eightfold explanation generation speedup on comparable or weaker hardware. An ablation study confirms that joint modeling is essential: removing node features reduces predictive faithfulness by 56% and generating features independently of structure reduces ground-truth faithfulness by 24%. This work demonstrates faithful, node-feature aware, model-level explainer for heterogenous GNNs is achievable from a single, efficient model.

Keywords: Explainable AI, flow matching, graph neural networks, heterogenous graphs, model-level explanation

Track 4: Medical and Health Informatics

Hybrid Ensemble of U-Net and Transformer-U-Net for Multi-class Brain Tumor Segmentation

¹Uthsara, B.A.D.E.*, ²Liyanaarachchi, K.L.P.P.
ethmal.20221590@iit.ac.lk, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

Brain tumor segmentation from Magnetic Resonance Imaging (MRI) scans remains a critical challenging task in medical image analysis, as tumors exhibit high variability in shape, size, and location. Existing deep learning approaches primarily rely on single-model architectures such as U-Net variants and different encoder decoder enhancements, which are limited in their ability to simultaneously capture local spatial features and global contextual dependencies. This limitation reduces model robustness and segmentation performance accuracy, particularly for multi-class tumor segmentation, under noise and MRI imaging variations.

This study proposes a novel hybrid ensemble framework by combining U-Net and Transformer U-Net architectures to address the identified limitations. The BraTS 2021 dataset was used for training, validation, and testing with a 70%, 15% and 15% split. Data augmentation techniques were applied to improve model generalization. Four ensemble strategies were designed, implemented, and compared: Soft-probability Voting, Weighted OR Voting, Attention-based Fusion and Pixel-wise Stacking. Models were evaluated using pixel-wise accuracy, Dice Coefficient, Intersection over Union (IoU), precision, sensitivity, specificity, and F1-score.

Within the proposed approaches, the Attention-based Fusion ensemble achieved the best overall performance with a Mean Dice Coefficient of 0.8774, Mean IoU of 0.7910, and pixel-wise accuracy of 0.9944. The Pixel-wise Stacking ensemble also delivered strong results with a Mean Dice Coefficient of 0.8758, Mean IoU of 0.7888, and accuracy of 0.9944. Both approaches outperformed the individual base models (U-Net: Dice Coefficient = 0.8627; Transformer U-Net: Dice Coefficient = 0.8650), demonstrating that adaptive ensemble fusion strategies significantly enhance multi-class brain tumor segmentation performance and robustness.

Keywords: Brain Tumor Segmentation, Ensemble Learning, MRI, Transformer U-Net, U-Net.

Leveraging Knowledge Distillation for Fine-Grained Non-Proliferative Diabetic Retinopathy Classification

¹Perera, G.*, ²Caucidheesan, K.

glenperera2002@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Diabetic Retinopathy is one of the main causes of vision loss worldwide. Early Detection is essential for preventing blindness. Diabetic Retinopathy consists of main two stages, which is Non Proliferative stage and the Proliferative stage. The early stages which consists of mild, moderate and severe stages belongs to the non-proliferative stage, which is commonly known as the early stages. Deep Learning models have demonstrated good performance in the categorization of general diabetic retinopathy but fail in fine-grained classification due to the high computational complexity and deployability issues, which involves differentiating between the mild, moderate, and severe sub-stages of non-proliferative diabetic retinopathy, which are clinically identical.

To address this challenge, a knowledge-distilled lightweight framework for fine-grained non-proliferative diabetic retinopathy classification was developed. A high capacity FastViT model is acting as the teacher model and the MobileViT-S acts as the student model. The distillation process ensures that the student model (MobileViT-S) retains the accuracy of the teacher (FastViT) while reducing the model complexity. The pipeline includes robust preprocessing of fundus images, data augmentation for class balancing.

The distilled MobileViT-S model has got an accuracy of 87.78% with the diabetic retinopathy dataset. This represents a reduction of more than 98.2% from the baseline for the FastViT teacher model. More importantly, this model achieves a highly competitive level of accuracy with 41.5% fewer model parameters and an average inference time of 14.68 ms, which is two times faster than the baseline. To validate the real world applicability, the lightweight model was successfully deployed and tested in low resource constrained devices, which features a real-time diagnostic web interface enabling clinicians to correct predictions and facilitate continuous model retraining.

Keywords: Diabetic Retinopathy NPDR Classification, Fast-ViT, Knowledge-Distillation, MobileViT-S

BrainXplain: Towards Explainable Brain Tumor Analysis - Segmentation and Uncertainty Aware Report Generation Pipeline

¹Aamil, A.A., ²Guhanathan, P.

aliyar.20211096@iit.ac.lk, ¹University of Westminster, United Kingdom

²Informatics Institute of Technology, Sri Lanka

Brain tumor diagnosis relies heavily on expert interpretation of multi-sequence 3D MRI, making it time-consuming and prone to inter-observer variability. Manual radiology reporting is further challenged by increasing imaging workloads and shortages of trained radiologists. This growing discrepancy between imaging volume and available expertise often leads to diagnostic delays, cognitive fatigue, and inconsistent reporting terminologies. Although deep learning has improved tumor segmentation and anomaly detection, current automated MRI report-generation systems remain clinically limited by weak evidence grounding, hallucination, and a lack of adaptability in healthcare settings due to uncalibrated uncertainty. To address this gap, this project proposes a trustworthy and adaptable brain tumor analysis and reporting pipeline that integrates segmentation, explainability, and evidence-grounded report generation.

The study employs a multimodal deep learning approach that integrates 3D semantic segmentation with a large vision-language model to generate uncertainty-aware radiology narratives. Publicly available datasets, including BraTS and related text-annotated resources were utilized with standardized preprocessing. A fine-tuned Swin UNETR model was deployed to segment the main tumor regions. Crucially, the pipeline integrates 3D Grad-CAM for visual explainability and Monte Carlo (MC) Dropout for voxel-wise uncertainty quantification. By specifically focusing on epistemic uncertainty and leaving aleatoric uncertainty for future evaluation, the framework establishes a foundation for a highly adaptable healthcare system. To bridge the visual and textual domains, quantitative radiomic features, extracted tumor volumes, atlas-based localization, and XAI artifacts are combined with these uncertainty measures into a unified evidence bundle. This bundle ultimately conditions a MedGemma-based model to produce structured, reliability-aware reporting.

The fine-tuned segmentation model achieved a Dice score of *0.8802* across the whole tumor, tumor core, and enhancing tumor regions, improving upon the pretrained baseline of *0.8513* while reducing the 95% Hausdorff Distance from *7.5545* mm to *6.3890* mm. Furthermore, an ablation study evaluating the impact of different pipeline components demonstrated that the full configuration yielded the most coherent and well-grounded reports. Language-model evaluation confirmed this, showing a mean Evidence Grounding and Claim Verification Rate of *0.816*, alongside a Claim Support Coverage of *0.867*. Advancing this research involves comprehensive evaluations with clinical experts to rigidly validate its diagnostic utility and further refine its adaptability for real-world clinical application. Ultimately, this framework represents a significant step forward in automated brain tumor analysis, bridging the gap toward trustworthy neuro-oncology workflows.

Keywords: Brain Tumor Segmentation, Explainable AI, Multimodal Deep Learning, Radiology Report Generation, Uncertainty Quantification

MQML-Inspired Hybrid Quantum-Classical Classification of TCGA-BRCA Tumor and Normal Samples Using Integrated Multi-Omics Features

¹Kodithuwakku, K.A.*, ²Poravi, G.

aweesha.20220538@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Breast cancer remains one of the most prevalent and molecularly heterogeneous malignancies globally, where accurate and early classification is critical for improving patient outcomes. Conventional machine learning approaches applied to multi-omics data encompassing RNA-Seq gene expression, miRNA expression, and DNA methylation face persistent challenges including extreme feature dimensionality, inter-modality heterogeneity, class imbalance, and susceptibility to overfitting. Meanwhile, quantum machine learning (QML) research has largely remained confined to synthetic or small-scale datasets, leaving its clinical applicability largely unvalidated.

This work presents QuantOmicBC, a hybrid quantum-classical machine learning framework for the classification of tumor and normal samples from the TCGA-BRCA multi-omics dataset. The pipeline employs a four-stage feature selection strategy combining Welch's t-test, Benjamini-Hochberg false discovery rate (BH-FDR) correction, Cohen's d effect size filtering, and Random Forest importance ranking to reduce over 485,000 raw features to a computationally tractable set. Selected features are amplitude-encoded into Quantum Neural Networks (QNNs) implemented using PennyLane integrated with PyTorch via the TorchLayer API. Two model variants are evaluated: MQML-16, a 4-qubit circuit with 14 trainable parameters, and MQML-64, a 6-qubit circuit with 38 parameters. SMOTE-based oversampling is applied exclusively within training folds to prevent data leakage, and Youden's J statistic is used to optimize classification thresholds. The framework is deployed as a FastAPI backend on Railway with a React 18 frontend on Vercel.

Experimental evaluation on a blind held-out test set demonstrates that MQML-64 achieves AUC-ROC = 0.9992, MCC = 0.9449, F1-score = 0.9960, and perfect tumor sensitivity of 1.000 with zero false negatives. MQML-16 achieves AUC-ROC = 0.9424 with only 14 trainable parameters, illustrating parameter efficiency as a meaningful form of quantum advantage. Quantum-weight sensitivity analysis, combined with Random Forest feature importance, enables biologically interpretable biomarker identification across omics modalities. These results demonstrate that compact quantum circuits can achieve clinically competitive performance on real-world multi-omics cancer data, bridging a meaningful gap between biomedical informatics and quantum computing.

Keywords: Biomarker Discovery, Breast Cancer Classification, Hybrid Quantum-Classical Model, Multi-Omics Integration, Quantum Machine Learning (QML)

Adaptive Label Refinement and Multi-Signal Robust Training for HAM10000 Skin Lesion Classification under Simulated Feature-Dependent Label Noise

¹Silva, S.*, ²Poravi, G.

savindi.20220375@iit.ac.lk, ¹University of Westminster, United Kingdom

²Informatics Institute of Technology, Sri Lanka

Deep learning models have demonstrated strong performance in skin lesion classification; however, their reliability is often reduced by noisy labels in medical datasets. In dermatology, annotation errors are commonly feature-dependent, where visually ambiguous lesions are more likely to be mislabeled due to diagnostic uncertainty and inter-expert disagreement. Such label noise can cause neural networks to memorise incorrect annotations, leading to reduced generalisation and unreliable predictions. Existing approaches frequently assume class-conditional noise or rely on loss-based sample filtering, which may discard valuable but difficult training examples and fail to adequately address feature-dependent noise.

This research proposes an adaptive label refinement and multi-signal robust training framework for skin lesion classification under simulated feature-dependent label noise. The framework introduces a two-stage SIEVE-and-REHAB strategy. The SIEVE stage combines multiple reliability signals, including loss dynamics, uncertainty estimation, prediction confidence, feature-space consistency, and prediction stability, to identify potentially noisy samples. The REHAB stage employs a teacher-student learning mechanism with adaptive pseudo-labelling to refine suspicious labels while preserving valuable hard examples. Unlike conventional filtering approaches, the framework aims to correct and reuse uncertain samples rather than simply removing them from training.

Experiments were conducted using the HAM10000 dermoscopic image dataset under simulated instance-dependent noise settings. A pretrained EfficientNet-B0 model was used as the backbone architecture and evaluated using Accuracy, Macro-F1 Score, and Matthews Correlation Coefficient (MCC). Results demonstrated improved robustness, training stability, and classification performance compared with baseline training under noisy conditions. Ablation studies further confirmed the contribution of both the SIEVE and REHAB components to overall framework effectiveness and highlighted the benefits of combining multiple reliability signals for noise estimation.

The findings indicate that adaptive label refinement and multi-signal reliability estimation can improve learning under feature-dependent label noise and support the development of more reliable deep learning systems for medical image classification. The proposed framework provides a practical and extensible approach for improving robustness in noisy medical datasets and may contribute to the development of more trustworthy computer-aided diagnostic systems.

Keywords: Deep Learning, Instance-Dependent Noise, Label Noise, Robust Training, Skin Lesion Classification

MindScape: An Explainable Game-Driven Behavioral AI Framework with Graph Neural Networks for Early Cognitive Risk Screening in Children

¹Sundareswaran, S. *, ²Athuraliya, B.

surruthisha.20222357@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Early cognitive decline is no longer a condition exclusive to aging populations. Among children of the current generation who have grown up entirely immersed in digital environments, early-stage cognitive-risk markers including shortened attention spans, weakened working memory and reduced executive control are becoming increasingly prevalent. Chronic exposure to hyper-stimulating content, autoplay media and dopamine-driven gaming platforms has been linked to structural brain changes including reduced grey-matter volume and impaired prefrontal function, collectively accelerating the onset of cognitive vulnerability at younger ages. Conventional instruments such as the Mini-Mental State Examination (MMSE) and Montreal Cognitive Assessment (MoCA) lack the sensitivity to capture subtle behavioral deterioration and are perceived as slow and anxiety-inducing by children accustomed to fast-paced interactive environments, rendering them both clinically insufficient and ecologically invalid for pediatric screening. There is a compelling need for a detection framework that is technically rigorous, developmentally appropriate and embedded within the digital contexts children already inhabit. This paper presents MindScape, an explainable game-driven behavioral AI framework that repurposes the digital environment as a clinical diagnostic instrument. Neurocognitive assessment tasks, validated against established clinical proxies, are embedded within a narrative-driven game to passively extract a high-resolution behavioral matrix capturing reaction latency, error patterns, memory span and cognitive flexibility without disrupting the child's experience. A novel Pediatric Feature-Space Calibration Layer maps raw gameplay telemetry across four cognitive domains, namely processing speed, cognitive flexibility, short-term memory and working memory, into adult-normative Neurocognitive Performance Test (NCPT) score ranges using developmentally grounded regression formulas. This enables zero-shot inference on pediatric profiles from a model trained on 318,000 clinical records without child-specific retraining. The predictive engine employs a Cognitive-Aware Graph Attention Network v2 (CA-GATv2) stacking ensemble combined with LightGBM, augmented by a Residual Informatics Bridge to preserve clinically meaningful subtest-level feature fidelity across deep graph neural network layers. Clinical interpretability is delivered through a post-hoc Explainable AI (XAI) layer based on counterfactual explanations, generating structured clinical reports with per-domain performance breakdowns and plain-language recommendations that communicate both the predicted risk classification and the behavioral adjustments that would alter a child's outcome. Evaluated on a held-out test set of 4,000 nodes, the CA-GATv2 ensemble achieved 85.6% accuracy, a Macro-F1 of 0.8612, AUC-ROC of 0.986 and a Quadratic Weighted Kappa of 0.9108, reflecting near-perfect clinical agreement by established psychometric standards. The at-risk class yielded a precision of 0.95 and recall of 0.91, the strongest per-class result, directly supporting reliable identification of children requiring clinical follow-up. An ethics-approved pilot study with 50 children aged 7 to 11 validated the calibration layer: 42 children (84%) were classified as high performance, 8 (16%) as average and 0 as at-risk, accurately reflecting the healthy school cohort. This work constitutes the first application of GNN-based multi-class classification to the NCPT battery and was published at

IEEE ICAIIC 2026, Tokyo. MindScape demonstrates that the same digital landscape implicated in early cognitive decline can be systematically reengineered as a scalable, non-invasive and explainable screening tool, offering clinicians, educators and caregivers a credible pathway to identifying at-risk children before difficulties become entrenched.

Keywords: behavioral biomarkers, cognitive risk screening, explainable AI, gamification, graph neural networks, HCI, Deep learning

Generative Augmentation of Anatomical-Functional Imaging for Rare Lung Cancer Subtype Classification

¹Payagalage, R.*, ²Yapa, P.

rumeth.20222449@iit.ac.lk, ¹Robert Gordon University, United Kingdom,

²Informatics Institute of Technology, Sri Lanka,

Lung cancer remains the leading cause of cancer related mortality worldwide, rendering early and accurate subtype classification essential for effective treatment and improved patient outcomes. Among its many subtypes, the rare ones present significant diagnostic challenges due to their high visual similarity with more common variants in medical imaging. Developing reliable Artificial Intelligence (AI) tools to address this challenge demands large volumes of high-quality annotated medical images. However, obtaining sufficient data for rare subtypes remains a significant limitation, given their low prevalence in clinical settings. Attempts to overcome this shortage through generative augmentation models often encounter a well known problem, mode collapse. This occurs when the model fails to capture the full diversity of the training data distribution and instead produces the dominant samples. This study aimed to develop and evaluate a generative image augmentation framework capable of synthesizing realistic and diverse medical scans for rare lung cancer subtypes, mitigating the mode collapse issue, and assessing the impact of these synthetic samples on downstream classification tasks.

A class-conditioned Denoising Diffusion Probabilistic Model was developed to synthesize Positron Emission Tomography (PET) scans, which capture metabolic activity within the body, for three rare lung cancer subtypes. As lung region annotations were unavailable, a threshold-based technique was applied to estimate lung regions. The fidelity, quality, and diversity of the generated images were evaluated using five metrics, including Structural Similarity Index Measure (SSIM) and Fréchet Inception Distance (FID). To assess the clinical impact of the synthetic data, a downstream classification framework was developed. Separate classifiers were trained on PET and CT imaging data and subsequently combined using late fusion techniques.

The best performing threshold configuration achieved SSIM of 0.827/1.0 and FID of 25.57 (0 is optimal) confirming the realism of the generated samples while mitigating inter-class mode collapse. Incorporating synthetic images into the training data improved overall classification accuracy by 9.46% to 12.81% compared to baseline models trained without augmentation. The probability-based late fusion of the best performing classifiers achieved a state-of-the-art accuracy of 98.88% across four lung cancer subtypes: Adenocarcinoma, Small Cell Carcinoma, Large Cell Carcinoma, and Squamous Cell Carcinoma. Domain expert evaluation with three independent medical professionals further confirmed the clinical realism of the generated images and the practical applicability of the proposed framework. These findings demonstrate that integrating generative image augmentation with PET-CT imaging is an effective and clinically relevant approach for classifying rare lung cancer subtypes. Ultimately, provides a practical pathway towards developing more accurate and generalized AI diagnostic tools in data limited medical imaging settings.

Keywords: Anatomical-Functional Imaging, Generative Augmentation, Mode Collapse, PET-CT Fusion, Rare Lung Cancer

FLEX-Med: A Federated Knowledge Distillation Framework for Acute Lymphoblastic Leukemia Cell Classification under Non-IID Data and Model Capacity Heterogeneity

¹Sathiyathan, S.D.*, ²Poravi, G.

shrenik.20210735@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Accurate diagnosis of Acute Lymphoblastic Leukemia (ALL) from peripheral blood smear images remains challenging due to reliance on manual microscopic analysis, which is time-consuming and subject to variability. While deep learning has shown strong potential for automated classification, its adoption across institutions is limited by restricted data sharing, non-IID data distributions, and differences in model capacities. This research proposes a federated learning framework for robust ALL classification under data and model heterogeneity.

To address this, FLEX-Med, a logit-based federated knowledge distillation framework, was developed to enable collaborative learning without parameter sharing. Using the C-NMC 2019 dataset, 12,529 labeled images were partitioned into local training data, a shared public anchor dataset, and a held-out test set. Heterogeneous EfficientNet-B0, B1, and B2 models were deployed across clients, while data heterogeneity was simulated using Dirichlet partitioning with a baseline $\alpha = 1.5$, alongside $\alpha = 0.5$ and 1.0 for evaluation. Knowledge exchange was achieved through shared logits on the anchor dataset, followed by local distillation.

Results demonstrate that the framework achieves stable and competitive performance under heterogeneous federated settings. For $\alpha = 1.5$, the system produced strong results across all model variants, outperforming or closely matching local baselines. Under more extreme heterogeneity ($\alpha = 0.5$), it maintained reliable performance with an average accuracy of 75.3% and F1-score of 78.8%, confirming its viability under non-IID data distributions and varying model capacities.

Keywords: Acute Lymphoblastic Leukemia, Data Heterogeneity, Knowledge Distillation Medical Image Classification, Model Heterogeneity,

DCTFusion: A Dual Branch CNN-Transformer Framework for Local-Global Feature Fusion in Diabetic Retinopathy Classification

¹Danishkar, S.*, ²Poravi, G.

danishkar.20210779@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Diabetic retinopathy (DR) is one of the leading causes of preventable blindness worldwide, making accurate severity grading essential for early diagnosis and timely clinical intervention. Existing automated DR classification systems are predominantly based on Convolutional Neural Networks (CNNs), which effectively identify localized retinal lesions such as microaneurysms, hemorrhages, and exudates. However, CNNs often struggle to capture long-range spatial dependencies and global contextual relationships across retinal regions that are important for distinguishing between adjacent DR severity stages. Transformer-based models address this limitation by modelling global context but generally require substantial computational resources and large annotated datasets. Furthermore, many existing hybrid CNN-Transformer approaches rely on static feature fusion strategies, resulting in limited interaction between local and global representations.

This research proposes DCTFusion, a dual-branch CNN-Transformer architecture designed to effectively integrate localized lesion information with global retinal context for five-stage DR severity classification. The framework employs EfficientNet-B4 as the CNN branch for local feature extraction and Swin Transformer as the transformer branch for capturing long-range dependencies. A cross-attention fusion module is introduced to enable adaptive bidirectional interaction between the two feature domains, enhancing complementary feature learning while reducing redundant representations. The model was trained using publicly available retinal fundus image datasets with preprocessing, augmentation, and class imbalance mitigation techniques to improve robustness and generalization.

Experimental evaluation on the APTOS 2019 dataset demonstrated that DCTFusion achieved a Quadratic Weighted Kappa (QWK) score of 0.8639 and an overall classification accuracy of 80.55% for multi-class DR grading. Compared with a standalone CNN baseline, the proposed framework improved QWK by 7.96% and accuracy by 5.10%. These findings indicate that adaptive local-global feature fusion through cross-attention can significantly enhance diabetic retinopathy severity classification and support the development of reliable computer-aided screening systems for clinical practice.

Keywords: Cross-attention fusion, Diabetic retinopathy, EfficientNet-B4, Medical image classification, Swin Transformer.

MURMUR: Explainable and Personalised Temporal Deep Learning for Breath-Based Augmentative Communication

¹Nawas, N.*, ²Athuraliya, B.

nafla.20220352@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Individuals with severe motor impairments, such as those affected by amyotrophic lateral sclerosis (ALS) and locked-in syndrome, frequently face profound communication challenges despite retaining voluntary respiratory control. Existing augmentative and alternative communication (AAC) systems remain inaccessible to many in this group due to their reliance on specialised hardware, limited personalisation capabilities, and a lack of transparent decision-making. MURMUR addresses these limitations by proposing a breath-based AAC system that classifies intentional short and long breath puffs captured through a standard consumer microphone.

The system employs a Multi-Scale Temporal Convolutional Network (MS-TCN) with four parallel dilated convolutional branches (dilation rates 1, 2, 4, and 8) operating on three-channel Mel-delta spectrograms extracted from raw respiratory audio. This architecture captures both fine-grained and long-range temporal dependencies in breath signals without requiring recurrent components. Training is conducted on the Coswara respiratory dataset (Bhattacharya et al., 2023) using a subject-wise 70/15/15 train-validation-test split across 2,542 subjects, ensuring that evaluation reflects cross-subject generalisation rather than within-subject overfitting. A Prototypical Network (ProtoNet) personalisation layer enables few-shot adaptation to individual users by constructing class prototypes from a small number of calibration samples, without requiring any retraining of the global model. Classification confidence is derived from temperature-scaled cosine similarity between query embeddings and stored prototypes. Explainability is supported through Grad-CAM visualisations of Mel-spectrogram regions, translated into plain-language acoustic descriptions for caregiver interpretation. Evaluation on 760 samples from 382 unseen subjects yields a global macro F1-score of 0.726 and a classification accuracy of 72.6%. Cross-subject ProtoNet testing demonstrates monotonic improvement from $F1 = 0.633$ at $k = 1$ support sample to $F1 = 0.730$ at $k = 10$, matching the fully trained global model without any retraining. Within-subject personalisation using augmentation-based calibration achieves $F1 = 0.984$, representing a 35.5% relative improvement over the unpersonalised baseline. An ablation study confirmed that removing Cepstral Mean and Variance Normalisation (CMVN) resolved systematic preprocessing inconsistencies between training and inference, and that a confidence threshold of 0.50 yields near-zero rejection rates appropriate for clinical deployment. MURMUR demonstrates that a lightweight, consumer-hardware-compatible breath-based AAC system can achieve strong classification performance while supporting user-level personalisation and caregiver-facing explainability within a single deployable architecture. Expert evaluation by an AAC specialist, applied scientists, and an individual with lived AAC experience confirmed the practical and clinical relevance of the system.

Keywords: assistive technology, breath gesture classification, explainable AI, few-shot personalisation, temporal convolutional network.

SymAD-ECNN: Geometry Aware Equivariant Learning for Anomaly Detection in Biomedical Imaging

¹Badurdeen, R., ²Wijesinghe, N.

riha.20220701@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri-Lanka

Conventional deep learning models for biomedical anomaly detection often suffer from ‘geometric blindness’, limiting their ability to handle spatial transformations such as rotations. This results in a reliance on computationally expensive data augmentation and large annotated datasets, which are scarce in clinical environments. Additionally, high false-positive rates and limited explainability reduce the trustworthiness and adoption of AI in brain MRI analysis.

This research proposes SymAD-ECNN, a geometry-aware unsupervised framework for anomaly detection in biomedical imaging, particularly focusing on brain tumours. This is based on an Equivariant Convolutional Autoencoder (ECNN-AE), which encodes rotational symmetry to ensure consistent feature representation without extensive augmentation and is trained exclusively on healthy T1-weighted MRI scans from the IXI dataset, under a one-class classification paradigm. The encoder is designed to be equivariant to geometric transformations, while the latent bottleneck enforces invariant representations, enabling robust modeling of normal anatomical structures. The decoder reconstructs the input image from this invariant representation, and anomalies are identified through reconstruction discrepancies.

The proposed model is validated using both IXI and BraTS (tumour) datasets, demonstrating strong anomaly detection performance, achieving an AUROC of 0.8625 and an AUPRC of 0.9197, outperforming standard CNN-based baselines. The model demonstrated high reconstruction fidelity, with a MSE of 0.0043 and a SSIM of 0.8725. Furthermore, reconstruction error maps support improved transparency. Collectively, these findings demonstrate that incorporating geometric priors directly into deep learning architectures yields a robust alternative to augmentation-heavy pipelines, offering a promising direction in this domain.

Keywords: Anomaly Detection, Biomedical Imaging, Equivariant Learning, One-Class Classification, Reconstruction Based Modeling

Optimizing Adult Glioma Segmentation via Geodesic-Guided Multi-Task Learning and Dynamic Loss Weighting

¹Sendanayake, J.D.*, ²Yapa, P.

janaka.20221496@iit.ac.lk, ¹Robert Gordon University (RGU), United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Gliomas are aggressive brain tumors that naturally grow into surrounding healthy tissue, making them difficult to delineate on MRI scans. While current deep learning models can estimate total tumor volume, they frequently fail to identify precise boundaries. Standard models rely on simple binary masks. A major flaw with this approach is the lack of spatial information in binary masks; a pixel at the fuzzy, uncertain boundary of a tumor is assigned the exact same label as a pixel deep inside the solid core. Because they ignore physical depth, algorithms struggle to learn the true shape of the tumor.

The primary objective of this research was to explore how Multi-Task Learning (MTL), which involves training models to do multiple related tasks at the same time, can be utilized alongside a dynamic loss function to significantly improve the accuracy of tumor boundary detection, demonstrating that a superior supervisory signal can compensate for architectural simplicity.

We developed a lightweight Dual-Decoder U-Net model that performs standard tumor region mapping alongside geodesic regression, which predicts the continuous physical distance to the tumor's boundary. To successfully balance these competing tasks without manual tuning, we implemented a dynamic loss function based on homoscedastic uncertainty. This approach automatically adjusts the task weights based on their relative difficulty during training. Additionally, a consistency rule ensures the predicted boundaries logically align with the predicted tumor regions.

The MTL model was tested on the BraTS 2023 dataset. Despite having only 9.5 million parameters, the framework achieved competitive performance in in-distribution data alongside much larger, state-of-the-art systems like the 31-million parameter nnU-Net. More importantly, the model demonstrated better performance in OOD datasets, including pediatric tumors, meningiomas, and scans from Sub-Saharan African populations. For example, on the African dataset, it reduced boundary errors by 55% compared to leading benchmarks and maintained structural stability where larger models experienced severe failures.

The findings demonstrate that combining MTL and geodesic regression with an uncertainty weighted dynamic loss function provides a highly effective alternative to simply increasing model size. This approach provides a robust and reliable framework for evaluating segmentation algorithms and advancing computational methods in medical image analysis.

Keywords: Dynamic loss function, Geodesic regression, Homoscedastic uncertainty, Multitask learning, Spatial information.

Track 5: Robotics, Security and Applied AI

AURA: Real-Time Environment Adaptation Using Neuro-Fuzzy Procedural Content Generation with Behavioural Clustering

¹Geevinda, K.W. J. P.*, ²Athuraliya, B.
pasindu.20212016@iit.ac.lk, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

Procedural content generation (PCG) in modern games is typically static or reactively unstable, failing to adapt to player behaviour across combat, collection, and exploration, resulting in reduced engagement, cold-start bias, and repetitive gameplay. AURA is a calibration-first, telemetry-driven framework designed for real-time adaptation.

Phase 1 established a neutral baseline involving 7 participants across 3 modes, utilizing a z-score NeutralityScore. Phase 2 modelled behaviour using K-Means soft memberships and a canonical ANFIS-inspired adaptation surface, approximated by a $6 \rightarrow 16 \rightarrow 8 \rightarrow 1$ MLP surrogate. Phase 3 deployed a decoupled TypeScript backend integrated with Unreal Engine, applying bounded PCG updates every 30 seconds.

The surrogate model achieved a Test R^2 of 0.9264 (95% CI [0.874, 0.954]) and an MAE of 0.0127. Clustering (3,240 windows) yielded a silhouette score of 0.3752. Sensitivity analysis identified `soft_combat` as the dominant feature (0.1449). A delta ablation showed that removing temporal features reduced the R^2 from 0.9726 to 0.5891 (-38.35 pp), confirming that soft membership alone is insufficient. Runtime latency was maintained at <200 ms, with 10 out of 12 requirements successfully validated and supported by 62 automated tests.

Keywords: Adaptive procedural content generation, Behavioural clustering, Neuro-fuzzy surrogate inference, Real-time game AI, Telemetry-driven adaptation.

Real-Time Vision-Based Human-Robot Interaction Framework for Low-Cost Embedded Robotic Arms

¹Dep, W.A.P.*, ²Embuldeniya, V.

anjana.20221307@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Industrial automation remains inaccessible to majority of small and medium-sized enterprises (SMEs) due to the high cost of robotic systems and the technical complexity of conventional programming methods. Traditional robotic arms require expert configuration using domain-specific languages or physical kinesthetic teaching, creating barriers that prevent widespread adoption. This research addresses these challenges by developing a vision-based robotic arm control system that eliminates the need for traditional programming by enabling the robotic arm to learn directly from human motion.

The proposed system employs computer vision and machine learning techniques to track human arm and hand movements in real time using a standard RGB camera. Google's MediaPipe library was employed for pose estimation capturing shoulder, elbow, and wrist joint coordinates and hand tracking to detect open/close gestures. A Kalman filter was applied to stabilize noisy angular data and reduce motion jitter. The calculated joint angles are transmitted wirelessly to an ESP32 microcontroller via a TCP socket, which drives servo motors on a custom-built four-degree-of-freedom robotic arm. SG90 servo motors were used for low-load joints, while MG90 metal-gearred servos handled higher-torque joints. Power stability was achieved through an external regulated supply, and migration from Arduino to ESP32 enabled fully wireless, OTA-updateable operation.

The system was evaluated through end-to-end latency characterization, jitter suppression analysis, functional testing, and benchmarking against inverse kinematics (IK)-based and learning-based alternatives. Experimental results confirm a mean end-to-end communication latency of 57.3 ms, with Kalman filtering achieving over 40% variance reduction across all joints and inter-frame angular variation below 4 degrees. The pipeline operates at approximately 10 frames per second on consumer CPU hardware, satisfying Nyquist sampling requirements for voluntary human arm movements. All ten functional use cases were satisfied, and the system maintained stable operation under induced Wi-Fi congestion without oscillatory instability. Benchmarking demonstrated clear advantages over existing paradigms in embedded feasibility and deployment simplicity.

The prototype demonstrates that deterministic joint-space retargeting, without inverse kinematics or GPU infrastructure, is sufficient for stable real-time imitation on commodity hardware. Future work will extend the system with closed-loop joint encoder feedback for improved positional accuracy, stereo vision to resolve depth ambiguity, and evaluation across broader motion types and user groups laying the groundwork for deployment-ready automation in resource-limited environments.

Keywords: computer vision, ESP32, Kalman filter, robotic arm, vision-based imitation learning.

GeMS: Application of Deep Learning Architecture for Multi-Agent Coordination in the Godot Engine

¹Mahanama, C.J.*, ²Wickremasinghe, B.
chanditha.20221469@iit.ac.lk, ¹University of Westminster, United Kingdom,
²Informatics Institute of Technology, Sri Lanka

The domain of video game AI, and specifically multi-agent AI, has historically been rigid and scripted, leading to predictable and deterministic NPC behaviour. Traditional scripted techniques such as Finite State Machines and Behaviour Trees offer lightweight computational overhead and intuitive design mechanisms, yet they fundamentally rely on predetermined logic and struggle to produce adaptive responses to dynamic environments. Solutions such as Multi-Agent Reinforcement Learning (MARL) have demonstrated the ability to produce emergent, adaptive coordination but these approaches are subject to astronomical computational costs and severe sample inefficiency, making them an ill fit for real-time game systems.

This research proposes the Graph Emergent MCTS-Inspired Strategy (GeMS) framework, a novel imitation learning pipeline designed to bridge offline strategic planning and real-time game inference. GeMS employs a Monte Carlo Tree Search (MCTS) oracle as a high-quality data generator, interacting with a custom Godot Engine environment to produce a rich dataset of optimal multi-agent tactical decisions across varied scenario configurations. A Graph Attention Network (GAT) is then trained on this dataset to learn the nuanced relational patterns and synergistic behaviours discovered by the MCTS. At runtime, the Godot game engine queries the deployed GAT model for real-time inference, effectively distilling deep strategic knowledge into computationally lightweight decisions for a squad of NPC agents.

The GeMS pipeline was implemented and evaluated through parametric testing of both the MCTS oracle and the GAT model, alongside system, functional, and non-functional testing. Optimal parameters identified include 50 simulations per state for the MCTS, and a GAT configuration of 200 epochs, hidden dimension of 64, 4 attention heads, learning rate of 0.001, and a batch size of 64. The fully trained and deployed model achieved a mean inference latency of 1.12 milliseconds and a maximum latency of 3.68 milliseconds, well within the sub-2ms non-functional target. Per-agent prediction accuracy relative to MCTS ground truth exceeded 70% across all agents in the squad, validating the framework's ability to faithfully distil MCTS tactical intelligence into a real-time deployable model.

GeMS contributes a generalizable, modular framework that decouples the strategic planning layer (MCTS), the learning layer (GAT), and the game inference layer (GDScript), making it accessible to indie and AA game developers without access to large-scale compute infrastructure. Expert evaluation confirmed the viability and novelty of the approach, with suggestions for future work including continual learning from live gameplay, humanization of agent behaviour, and scaling to larger and more diverse agent configurations.

A Hybrid Deep Learning Approach for Phishing Website Detection Using Transport and Application Layer Features

¹Aslam, A.A.*, ²Goonathilaka, O.

aathiq.20210466@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Phishing websites remain a major cybersecurity threat by exploiting user trust and bypassing traditional detection mechanisms through evolving obfuscation techniques. Existing approaches rely on either application-layer features, such as URLs and HTML content, or transport-layer characteristics like TLS metadata. However, these single-layer methods are increasingly ineffective, especially with the adoption of TLS 1.3, which limits visibility into encrypted traffic.

To address this, this research proposes a hybrid two-layer phishing detection model that integrates transport-layer and application-layer analysis using deep learning. The first layer analyzes TLS handshake and network-level features, while the second layer evaluates URL structures, HTML content, and DOM features. A combined decision mechanism is used to improve detection accuracy and reduce false negatives.

The model was evaluated using both custom-built and publicly available datasets with optimized preprocessing and feature engineering. It achieved an accuracy of ~0.94, precision of ~0.93, recall of ~0.97, and an F1-score of ~0.95. Notably, the false negative rate (FNR) was reduced to ~0.02-0.03, significantly improving detection reliability. Compared to baseline models, the proposed approach achieved a 4-6% performance improvement, particularly in minimizing missed phishing attacks.

Keywords: Cybersecurity, Deep Learning, Phishing Detection, TLS Features, URL Analysis

A Mesh-Free, Turbulence-Aware Physics-Informed Neural Network with Explainable AI for Airfoil Optimization at High Reynolds Numbers

¹Vidanya, H.V.G.H.*, ²Priyanayana, K.S.,
hasmaleevidanya@gmail.com, ¹University of Westminster, United Kingdom
²Informatics Institute of Technology, Sri Lanka

Aerodynamic shape optimization (ASO) of airfoils at high Reynolds numbers remains a challenging problem due to turbulence effects and the high computational cost of conventional Computational Fluid Dynamics (CFD) approaches. Traditional mesh-based solvers require repeated mesh generation and re-meshing during optimization, making iterative design expensive and time-consuming. Furthermore, most existing Physics-Informed Neural Network (PINN) frameworks are limited to low-Reynolds-number flows and provide limited interpretability of their predictions.

This research presents a mesh-free PINN framework for high-Reynolds-number airfoil analysis and optimization that integrates the Reynolds-averaged Navier-Stokes (RANS) equations with the Spalart-Allmaras one-equation turbulence model directly into the PINN loss function. Airfoil geometries are represented using PARSEC parameterization, enabling smooth and physically meaningful shape variations through a compact set of geometric design variables. A hybrid physics-guided learning framework is introduced, combining adaptive residual-guided refinement with multi-zone collocation sampling, including Sobol quasi-random interior points, near-wall Gaussian points, wake-region points, and airfoil-surface sampling to improve physics enforcement throughout the computational domain. Using automatic differentiation and the L-BFGS optimizer, the framework performs efficient gradient-based optimization to maximize the lift-to-drag ratio without requiring adjoint solvers.

To improve transparency and trustworthiness, four Explainable Artificial Intelligence (XAI) modules are incorporated to evaluate physical-law satisfaction, identify influential geometric parameters, visualize performance sensitivity, and highlight critical flow regions. The optimized two-dimensional airfoil design is further extended into a three-dimensional wing visualization to provide a more realistic representation of aerodynamic improvements. Experimental evaluation demonstrated a 90% improvement rate in lift-to-drag ratio optimization relative to the baseline PINN predictions when benchmarked against CFD data. The optimized airfoils consistently exhibited positive lift-to-drag ratio changes through increased lift and reduced drag, while the framework maintained stable training and practical computational performance through checkpoint-based learning.

By combining physics-based optimization, turbulence modeling, explainability, adaptive sampling, and visualization within a single end-to-end framework, this work advances practical PINN-based aerodynamic design for high-Reynolds-number applications.

Keywords: Aerodynamic Shape Optimization (ASO), High Reynolds Number, Physics-Informed Neural Networks (PINNs), Spalart-Allmaras (SA) Model, Turbulence Modeling.

AIM-IPS: A Five-Layer Adaptive Intrusion Prevention Framework with Cross-Pipeline Threat Amplification

¹Prabashwara, N.*, ²Udagedara, C.

nipuna.20221148@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Existing intrusion prevention systems enforce a binary block-or-allow decision, creating an unavoidable tradeoff between false positives and missed attacks, and operate network-layer and application-layer detection as independent silos leaving coordinated multi-stage attack campaigns undetected when each stage falls below the block threshold individually. This paper presents AIM-IPS, a production-deployed intrusion prevention middleware that addresses both limitations.

The central contribution is a cross-pipeline temporal correlation engine that shares threat state between the network-layer and application-layer detection pipelines in real time, amplifying fused threat scores by $1.3\times$ - $1.7\times$ when the same source IP triggers both network-layer reconnaissance and application-layer exploitation within a configurable time window. This mechanism correctly escalates 94.3% of coordinated kill-chain campaigns to block, with zero false escalations on clean traffic. A five-tier graduated response engine maps continuous fused threat scores to proportional actions (allow, delay, throttle, captcha, block), reducing legitimate request blocking from 7.1% to 1.9% while cutting automated attacker scan rates by approximately 99% at the throttle tier. An autonomous AI reasoning agent performs asynchronous forensic analysis with MITRE ATT&CK mapping and autonomous policy induction, processing 1,247 threat events and maintaining 144 validated adaptive detection rules in production.

Evaluation spans three levels of increasing ecological validity: benchmark performance on CICIDS2017; real-world validation on 1,000 production flows, maintaining 98.60% accuracy without retraining; and cross-dataset evaluation on UNSW-NB15, achieving 86.48% accuracy and 96.22% ROC-AUC under 5-fold cross-validation, with the unsupervised zero-day ensemble detecting 100% of novel attack categories. Combined application-layer evaluation on a 50,000-sample labelled dataset demonstrates 97% precision and 98% recall at a median end-to-end latency of 4.9 ms, confirming the feasibility of real-time inline deployment on commodity hardware.

Keywords: Cross-Pipeline Threat Correlation, Graduated Response, Intrusion Prevention System, Multi-Stage Attack Detection, Web Application Security

Plantar Pressure Heatmap Prediction from Monocular RGB Gait Video using BiLSTM with UNet and Heuristic-based Orthotic Recommendation

¹Palliyeguru, T.M.*, ²Poravi, G.

thusaraka@yahoo.co.uk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Custom orthotic prescription relies on force plates, pressure mats, and clinical assessment workflows that most patients cannot access, particularly in low-resource healthcare settings where such equipment is scarce or prohibitively expensive. Monocular video offers a cheaper alternative, but no prior work has demonstrated that plantar pressure distributions for baropodometric analysis and insole height adjustments can be recovered from a single RGB camera without pressure sensors at inference. This research gap motivates the present work, which asks whether such a pipeline is learnable and reports a proof of concept. VORMA extracts 56-dimensional biomechanical features from RTMPose-m keypoints estimated across 9-frame gait windows, feeds them into a two-layer bidirectional LSTM with temporal attention, and decodes bilateral plantar pressure heatmaps through a UNet-style decoder. The temporal attention mechanism allows the model to weight the most biomechanically informative frames within each gait window, improving its sensitivity to subtle pressure-shifting events during the stance phase. A deterministic rule-based engine then converts the predicted heatmaps into region-specific insole height recommendations, with MC Dropout uncertainty bands and clinician-verification flags. The full pipeline requires only a standard video camera and no pressure sensor is needed after training.

Evaluated on the MMVP dataset under leave-one-subject-out cross-validation across all 11 subjects, the model achieves a mean test RMSE of 2.113 ± 0.118 and SSIM of 0.503 ± 0.019 (95% CI RMSE: 2.043-2.183). It outperforms the Transformer baseline on all seven test metrics with 1.56x fewer parameters (0.393M vs 0.614M), a difference confirmed as statistically significant by paired t-test ($p < 0.01$). Three-seed reproducibility runs yield a standard deviation of 0.007 RMSE, confirming the results are not a training artefact. The orthotic engine reduces the Pressure-Weighted Risk Index by a mean of 37.6% across the test population ($t = 39.7$, $p < 0.001$), validated through a closed-loop surrogate simulation independent of any literature reduction curve. Orthotic technicians reviewed the prescription outputs and confirmed that the pressure-to-height mappings were consistent with standard clinical reasoning.

The results establish that a keypoints-only temporal model outperforms image-based approaches at this dataset scale, where convolutional encoders learn appearance shortcuts rather than biomechanical patterns. The system does not replace clinical judgment, but provides a low-cost, accessible, and interpretable starting point for camera-based orthotic screening. By demonstrating that meaningful pressure information can be inferred from ordinary video footage, this work lays groundwork for future clinical translation studies involving real patient populations and on-device deployment. More broadly, it positions camera-based gait analysis as a practical bridge between research-grade biomechanical instrumentation and everyday clinical practice.

Keywords: gait analysis, gait feature extraction, personalised orthotics, spatiotemporal gait modelling, video-based biomechanics

AeroSynth: A Physics-Informed Multi-Resolution Graph Framework for Aircraft Fuel Prediction and Route Optimization

¹Mohamed, A.F.*, ²Jawaadh, S.

aaysha.20221493@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The aviation industry faces continuous challenges in balancing operational efficiency with accurate fuel prediction and route planning, largely due to a critical lack of publicly accessible, tail-specific performance models. Operators are often forced to rely on generic manufacturer specifications that ignore individual engine degradation over time, while traditional flight planning software depends on rigid airways that struggle to dynamically bypass volatile weather. Furthermore, alternative deep learning prediction methods are plagued by opacity and a lack of physical constraints, preventing the critical trust required for flight-deck certification.

To overcome these barriers, this project introduces AeroSynth, a physics-informed machine learning framework integrated into a real-time route optimization system that fuses data-driven graph intelligence with core aerodynamic constraints. The implemented solution features a novel Physics-Informed Multi-Resolution Sensor Graph Kolmogorov-Arnold Network (PI-MRSG-KAN) trained on extensive Quick Access Recorder (QAR) data. Utilizing a "residual-corrector" loop, the model maps 25 distinct sensor streams into an adaptive graph network, blending physical principles with wavelets and learnable spline activation functions to capture transient sensor dependencies across all flight phases. This predictive engine is integrated into an interactive decision-support web application that continuously streams live operational parameters including dynamic weather feeds, winds, active FAA NOTAM alerts, and restricted airspace boundaries to evaluate and contrast optimized direct, layover, and multi-stop flight tracks.

Rigorous empirical testing on 410 held-out, flight-level test logs demonstrates that the PI-MRSG-KAN framework achieves exceptional predictive fidelity, securing a coefficient of determination (R^2) of 0.9814 and a Symmetric Mean Absolute Percentage Error (SMAPE) of 10.63%. Crucially, ablation testing reveals that anchoring the system with a physics-causal graph topology drops the Mean Absolute Error (MAE) by 44% compared to standard non-graph flat feature vector benchmarks. Furthermore, qualitative evaluation by domain experts confirms that the application effectively delivers clear, context-specific explanations of its route updates, successfully reducing pilot planning workloads. This integration validates that physics-informed learning can be operationalized into a practical, interpretable, and open-source tool, bridging the gap between high-fidelity prediction and real-world usability for aviation stakeholders.

Keywords: Aircraft Performance Modeling, Fuel Efficiency, Machine Learning, Route Optimization, Graph Neural Networks, Kolmogorov-Arnold Networks (KAN).

Per-Partner OCC Emotional Appraisal and Dual-Process Gating for Sustained Cooperation in Multi-Agent Reinforcement Learning

¹Ranasinghe, V.*, ²Peiris, S.

venukaranasinghe@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institution of Technology, Sri Lanka

Sustained cooperation in multi-agent reinforcement learning (MARL) remains an open challenge in social dilemmas such as the iterated Prisoner's Dilemma. Payoff-driven agents form cooperative clusters that are fragile: scalar reward signals carry no information about which neighbour defected or whether reciprocity is emerging. Humans stabilise cooperation by maintaining per-partner expectations and updating emotional states in response to violations. Bridging this gap between affective human reasoning and payoff-driven agents motivates this work.

We present E-MARL, equipping each decentralised agent with a per-partner OCC (Ortony-Clore-Collins) appraisal module tracking trust, anger, and guilt. Trust rises on positive expectation violations and decays over time; anger scales with trust invested upon defection; guilt accumulates when an agent exploits a cooperating partner and dissipates through prosocial actions. These dyadic states feed a dual-process emotion-gated Q-network, where appraisal vectors modulate the environmental hidden stream multiplicatively—so affect reshapes attended cues rather than merely augmenting the input.

E-MARL is validated on the 9×9 toroidal lattice Prisoner's Dilemma of Ren and Zeng (IJCAI 2024), with 81 independent DQN agents trained for 20,000 steps against a payoff-only baseline, across 16 and 13 seeds, using Welch t-tests and Cohen's d.

E-MARL raises terminal cooperation from 81.3% to 90.6% ($p = 0.005$, $d = 1.19$), converges 41% faster, and more than doubles stability. Social welfare improves 6.9% while the Gini coefficient falls 59% ($0.180 \rightarrow 0.074$), confirming emotional agents achieve higher collective outcomes and fairer distribution. Appraisal dynamics—trust rising, anger peaking then falling, guilt dissipating—align with OCC theoretical predictions.

Contributions include: (i) a per-partner OCC module with temporal decay and guilt-cooling; (ii) a dual-process emotion-gated Q-network in which appraisal gates modulate environmental representations; (iii) appraisal-based reward shaping for training acceleration; and (iv) a multi-seed empirical study spanning cooperation, stability, welfare, and fairness. These results establish cognitively grounded emotional appraisal as a strong inductive bias for MARL cooperation, with implications for autonomous vehicle coordination, federated learning, and distributed resource management.

Keywords: Dual-Process Theory, Emotion-Gated Networks, Multi-Agent Reinforcement Learning, OCC Appraisal, Social Dilemma

IntelF1: A Machine Learning Approach to Formula 1 Race Winner rediction and Evaluating its Betting Market Efficiency

¹Hapukotuwa, A.S.*, ²Ranawaka, G.

anuja.20222102@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Predicting the sole winner of a Formula 1 Grand Prix is much harder than predicting a podium finish, yet it is the specific outcome that matters most when placing a bet. Most existing machine learning work on F1 focuses on broader targets like top three finishes, points scoring positions, or seasonal standings, leaving a clear gap in models that estimate the probability of a single driver winning a race. This paper presents IntelF1, a machine learning framework that estimates per driver win probabilities and maps them to structured bet sizes using a modified fractional Kelly staking strategy. Three gradient boosted classifiers, XGBoost, CatBoost, and LightGBM, were trained and compared on Formula 1 race data spanning 2010 to 2024. The dataset was restricted to this period for two reasons, the blanket ban on mid-race refuelling in 2010 changed race strategy fundamentally, and the switch to V6 turbo-hybrid engines in 2014 made pre hybrid pace data non comparable to modern figures. A feature pipeline of 27 variables was constructed using rolling metrics rather than static career averages, including driver form, constructor reliability, and driver constructor synergy, each recalculated before every race. To handle the severe class imbalance inherent in F1 datasets, where exactly one of twenty drivers wins each race, cost-sensitive weighting was applied via the `scale_pos_weight` parameter, raising Top-1 classification accuracy from 49% to a range of 65-71% depending on the model. Economic performance was assessed across two evaluation windows. The first covered the final five races of the 2021 Formula 1 season as an initial stress test. The second used the complete 24 race 2024 season as a fully out-of-sample validation. The 2024 season produced seven different race winners from six constructors, the highest winner diversity in the modern era, making it a particularly demanding test for any prediction system. Across both windows, the IntelF1 non-linear staking strategy, which raises predicted win probabilities to the power of 1.9 before computing the Kelly fraction, outperformed both a market-favourite baseline and standard Kelly Criterion. The 2021 window produced a +70.2% aggregate ROI compared to -5.3% for the market favourite and +19.3% for standard Kelly. Across the full 2024 season the framework returned +171.50%, compared to -8.80% and +24.10% respectively. The largest returns occurred in races where bookmaker pricing anchored to driver reputation rather than circuit specific pace data, creating consistent mispricing that the model was able to identify. The framework has clear limitations. It operates entirely on pre-race features and cannot account for first lap incidents, mechanical failures, or safety cars that alter race outcomes. The non-linear staking function amplifies losses in high confidence incorrect predictions, as seen in races such as Miami and Austria in 2024. The reported ROI figures also assume ideal betting conditions with no stake limits or account restrictions, which are not guaranteed in live sportsbook environments. The upcoming 2026 FIA regulatory changes represent a structural break that will require feature re-weighting as new data becomes available.

Keywords: Betting Optimization, Capital Allocation, Formula 1, Gradient Boosting, Machine Learning

DeepEncryptNet: Privacy-Preserving Encrypted Traffic Monitoring via Metadata, Confidence Calibration, and Reinforcement Learning

¹Costa, P.*, ²Ramanayake, A.

praneeth.20211404@iit.ac.lk, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

The widespread adoption of encrypted communication protects user privacy but reduces the effectiveness of conventional network monitoring techniques that depend on payload inspection. This research presents DeepEncryptNet, a lightweight, privacy-preserving, and trust-aware framework for monitoring encrypted network traffic using only observable packet and flow metadata. The framework provides category-level traffic visibility without decrypting packets or accessing their contents, making it suitable for privacy-sensitive and resource-constrained monitoring environments.

DeepEncryptNet processes network flows and extracts statistical characteristics including packet sizes, packet timing, direction, packet counts, and flow-level distributions. A compact one-dimensional convolutional neural network with an attention mechanism learns local feature patterns and relationships among flow statistics. The backbone generates traffic-class predictions and a 128-dimensional embedding for subsequent analysis. Temperature scaling is selectively applied to improve the reliability of prediction confidence. A reinforcement learning decision agent then examines the predicted class probabilities together with maximum confidence, predictive entropy, and the margin between the two most probable classes. Based on these uncertainty signals, the agent either accepts the predicted traffic category or routes the flow to an unknown-handling path instead of forcing a potentially unreliable classification. Routed embeddings can subsequently be grouped for analyst review, future labelling, and the identification of previously unseen applications or potential tunnelling behavior.

The framework was evaluated using the ISCX VPN-NonVPN and HTTPS Traffic Classification datasets. The statistical-flow model achieved 86.98% accuracy and an 83.93% macro-F1 score on the ISCX dataset, while achieving 97.57% accuracy and a 96.33% macro-F1 score on the HTTPS dataset. Compared with rigid confidence-threshold rejection, the reinforcement learning policy preserved substantially greater usable coverage. It accepted 88.91% of ISCX samples and 99.44% of HTTPS samples while maintaining accepted-set macro-F1 scores of 90.79% and 96.75%, respectively. These findings demonstrate that separating traffic prediction from decision trust enables safer and more practical encrypted traffic monitoring.

Keywords: Confidence calibration; Encrypted traffic classification; Metadata-based monitoring; Reinforcement learning; Uncertainty routing

Proximal Policy Optimization with Weighted Priority Architectures for Resilient Spatial Navigation and Autonomous Disposal Routing in Mobile Edge Robotics

¹Abeygunasekara, S.*, ²Priyanayana, S.

sithilsandinu123@gmail.com, ¹University of Westminster, United Kingdom,

²Informatics Institute of Technology, Sri Lanka

Traditional waste management relies on fixed collection schedules and stationary receptacles, frequently resulting in overflow in high-traffic areas and inefficient manual labour. To transition urban hygiene toward a proactive model, this research presents an autonomous disposal routing framework utilizing Deep Reinforcement Learning (DRL) for a mobile holonomic waste receptacle. Upon detecting maximum internal capacity, the system initiates a spatial navigation policy to autonomously locate and dock with a centralized disposal facility. Designed for efficient execution on edge hardware, the Proximal Policy Optimization (PPO) agent is trained via a four-phase curriculum learning strategy. To navigate dynamic environments, the agent processes an 86-dimensional spatial array. This state space incorporates 72 normalized virtual LiDAR rays (providing 360-degree coverage up to 6 meters), a 2D relative goal vector, a 3D kinematic self-state V_x, V_y, ω , and a 9D dynamic obstacle tracking module. By continuously evaluating the distance, angle, and velocity of the three nearest pedestrians, this sub-array drives a custom 60/25/15 weighted priority tracking scheme. This ensures comprehensive spatial awareness, allowing the agent to execute complex evasive manoeuvres and navigate narrow corridors without inadvertently reversing into secondary flanking obstacles. Furthermore, the framework resolves the 'freezing robot' problem which is a common failure where the robot stops completely because the penalty for nearing an obstacle outweighs the reward for reaching the goal. To overcome this, the system dynamically reduces obstacle penalties by 40% whenever the robot actively moves toward its destination, encouraging it to confidently push through tight spaces rather than halting defensively.

The architecture was evaluated within a high-fidelity PyBullet simulation, utilizing extreme domain randomization to generate unpredictable office layouts populated by up to five dynamic pedestrians. The proposed model demonstrated robust performance, achieving an 80.0% navigation success rate. This significantly outperformed traditional A*+DWA mathematical planners (8.0% success) and deterministic TD3 baselines (42.0% success, primarily failing due to defensive freezing). An ablation study further validated the efficacy of the custom reward architecture: an ablated PPO trained without the 60/25/15 priority penalty prioritized shortest paths over safe clearances, resulting in a reckless 30.0% collision rate and a reduced 70.0% success rate. In contrast, the proposed PPO safely mitigated these collisions to 18.0%. Ultimately, these results substantiate that custom reward shaping and weighted multi-threat tracking can effectively guide holonomic edge robots through highly unstructured spaces, advancing the viability of automated dynamic logistics.

Keywords: Autonomous Navigation, Edge Computing, Obstacle Avoidance, Proximal Policy Optimization, Smart Waste Management



INFORMATICS
INSTITUTE OF
TECHNOLOGY



iciIT 2026
SYMPOSIUM

iR RESEARCH AND
INNOVATION UNIT

ISSN 3093-5202



9 773093 520007 >