ICTer 2026

26th International Conference on Advances in ICT for Emerging Regions

Book of Abstracts

04 - 05 November 2026

www.icter.lk

University of Colombo
School of Computing
UCSC



Book of Abstracts

26th INTERNATIONAL CONFERENCE ON ADVANCES IN ICT FOR EMERGING REGIONS



Disclaimer

No responsibility is assumed by the University of Colombo School of Computing for statements and opinions expressed by the contributors to this book of abstracts.

Abstracts of 26th ICTer International Conference (ICTer 2026)

Conference Website:https://icter.lk/

Copyright and Reprint Permission:
The copyright covers the exclusive right to reproduce and distribute the abstracts including reprints, photographic reproductions, microforms or any other reproductions of similar nature and translation. No part of the Book of Abstracts may be reproduced without the permission in writing from the copyright holder.
All rights reserved. Copyright ©2026, University of Colombo School of Computing.

MESSAGE FROM THE DIRECTOR



Dr. Ajantha Atukorale

Director
University of Colombo School of Computing
University of Colombo
Sri Lanka

It is my great pleasure, as Director of the University of Colombo School of Computing (UCSC), to warmly welcome all participants, distinguished researchers, academics, industry professionals, and emerging scholars to the International Conference on Advances in ICT for Emerging Regions (ICTer 2026).

ICTer represents a proud and enduring legacy in Sri Lanka's Computing research landscape. Its roots trace back to the International Information Technology Conference (IITC), a pioneering initiative that began in 1998. Building on this foundation, ICTer has continued to bring together diverse perspectives, foster intellectual exchange, and advance the role of Information and Communication Technology in addressing the challenges and opportunities of emerging regions. This legacy reflects UCSC's longstanding commitment to excellence in Computing education, research, and innovation.

We are living in an era in which Computing and digital technologies are transforming economies, reshaping societies, and redefining how we live, work, and interact. Yet, the true value of technological advancement lies not merely in what we can create, but in how effectively we apply it to solve meaningful problems, expand opportunities, and improve lives. In this context, ICTer provides a vital platform for exchanging ideas, presenting original research, challenging established thinking, and exploring innovative solutions to both regional and global challenges.

One of the enduring strengths of ICTer is its ability to bring together established researchers and the next generation of Computing scholars. By encouraging dialogue between academia, industry, and the wider research community, the conference creates opportunities for new perspectives, interdisciplinary collaboration, and lasting partnerships. I am particularly encouraged by the participation of emerging researchers, whose curiosity, creativity, and determination will help shape the future of Computing.

The conference also seeks to extend the reach and impact of high-quality research beyond the event itself. Selected papers will be considered for publication in the prestigious Springer Communications in Computer and Information Science (CCIS) series and the ICTer Journal published by UCSC, subject to the respective publication processes and requirements. These avenues provide valuable opportunities for researchers to disseminate

their findings, engage with the international research community, and contribute to the growing body of knowledge in Computing.

As we gather for ICTer 2026, I encourage every participant to engage openly, question thoughtfully, share generously, and collaborate enthusiastically. May the discussions initiated here lead to new discoveries, the connections made here grow into meaningful partnerships, and the ideas presented here inspire solutions that extend far beyond the conference.

On behalf of UCSC, I extend my sincere appreciation to the organising committee, reviewers, authors, partners, and everyone whose dedication has made ICTer 2026 possible. Your collective commitment sustains this important research tradition and strengthens our shared vision for the future of computing.

Let ICTer 2026 be more than a meeting of minds; let it be a catalyst for ideas that matter, innovations that empower, and collaborations that endure. Together, let us advance the frontiers of Computing and harness technology to build a future that is more intelligent, inclusive, connected, and beneficial to humanity.

I wish you all a stimulating, productive, and memorable ICTer 2026.

MESSAGE FROM THE CONFERENCE CO-CHAIRS



Dr. Shiromi Arunatileka,
Co-Chair, ICTer 2026
University of Colombo School of Computing (UCSC)
Sri Lanka



Dr. Roshan Rajapakse,
Co-Chair, ICTer 2026
University of Colombo School of Computing (UCSC)
Sri Lanka

It is our great pleasure to welcome you to the 26th International Conference on Advances in ICT for Emerging Regions (ICTer 2026). For more than a quarter of a century, beginning with the International Information Technology Conference (IITC) in 1998, this conference has brought together researchers, practitioners and innovators who share a common purpose: to build ICT solutions that are scalable, sustainable and meaningful for the people and communities of emerging regions.

This year's response from the research community was exceptional, with 332 submissions received. Following a rigorous review process, 40 papers were accepted for the Main Research Track along with 4 papers for the Industry R&D Track. This overall acceptance rate of about 13% reflects both the strength of the submissions and our commitment to upholding the high standards for which ICTer is known. Selected papers will be considered for publication in the Springer Communications in Computer and Information Science (CCIS) series and in the ICTer Journal.

A highlight of this year's program is the high-level panel discussion on the conference theme, "Artificial Intelligence (AI) and the Socio-Economic Future: Challenges and Opportunities." AI is reshaping economies, labor markets, public services and everyday life at remarkable speed. For emerging regions, this transformation brings both great promise and challenges. Bringing together distinguished experts from these communities, the panel will explore how nations like ours can utilize AI responsibly to build a more prosperous and equitable future.

A conference of this scale is made possible only through the dedication and commitment of many people. We extend our sincere thanks to all the authors who chose ICTer as the venue to share their work, and to the members of the Technical Program Committee and the reviewers who so generously contributed their time and expertise. We are also deeply

grateful to our international collaborators for their continued support, and to UCSC staff and student volunteers, whose hard work behind the scenes made this event possible.

Finally, we thank you, our participants, for being part of ICTer 2026. We hope you will find the sessions inspiring, the discussions stimulating, and the connections you make here lasting. Welcome to Colombo, and we wish you a rewarding conference.

ORGANIZING COMMITTEE

Co-Chairs of the Conference

- Dr. Shiromi Arunatileka, University of Colombo School of Computing
- Dr. Roshan Rajapakse, University of Colombo School of Computing

Main Organizing Committee

- Dr. Manju Wickramasinghe, University of Colombo School of Computing
- Dr. Randil Pushpananda, University of Colombo School of Computing
- Dr. Dinuni Fernando, University of Colombo School of Computing
- Dr. Tharindu Wijethilake, University of Colombo School of Computing

Conference Secretariat

- Ms. Madhumali Dhanawardana, University of Colombo School of Computing
- Ms. Venuka Fernando, University of Colombo School of Computing

Conference Administrative

- Ms. Deepani Chandralatha,University of Colombo School of Computing
- Mr. E.M. Gunarathna,University of Colombo School of Computing
- Mr. Kolitha Jayarathna,University of Colombo School of Computing
- Ms. Samudra Ekanayake,University of Colombo School of Computing
- Ms. Nuwanthi Jayasinghe,University of Colombo School of Computing

Advisory Committee

- Dr. Ajantha Atukorale, University of Colombo School of Computing
- Prof. Kasun De Zoysa, University of Colombo School of Computing
- Prof. Thiemo Voigt, Professor, Uppsala University, Sweden

Technical Program Committee Chairs

- A/Prof. Anuja T. Dharmaratne, Monash University, Australia
- Dr. Yohani S. Ranasinghe, University of Colombo School of Computing
- Dr. Tharindu Wijethilake, University of Colombo School of Computing

Honorary Advisory Committee

- Prof. Ryohei Nakatsu, Adjunct Professor, Kyoto University
- Prof. Raveendran Paramesran, Universiti Malaya

Publication Chair

- Dr. Thilini Lakshika, University of Colombo School of Computing

Technical Program Committee Co-Secretaries

- Mr. Jaliya Amarathunga, University of Colombo School of Computing
- Ms. Samudi Perera, University of Colombo School of Computing

Publication Sub Committee

- Ms. Hamsavasini Senthil, University of Colombo School of Computing
- Ms. Nadeesha Epa, University of Colombo School of Computing
- Mr. Prasad Wilagama, University of Colombo School of Computing
- Ms. Dinu Jayathissa, University of Colombo School of Computing
- Ms. Dilki Sewwandi, University of Colombo School of Computing
- Mr. Dilum Gamage, University of Colombo School of Computing
- Ms. Chethana Kawmini, University of Colombo School of Computing

Proceedings Editorial Board

- Dr. Yohani S. Ranasinghe,University of Colombo School of Computing
- Dr. Tharindu Wijethilake,University of Colombo School of Computing
- A/Prof. Suranga Nanayakkara, National University of Singapore
- A/Prof. Anuja T. Dharmaratne, Monash University
- Dr. Roshan Rajapakse, University of Colombo School of Computing
- Dr. Shiromi Arunatileka, University of Colombo School of Computing

Keynotes Speakers' Handling Committee

- Ms. Gayani Rupasinghe (Co-Chair), University of Colombo School of Computing
- Mr. Amod Pathirana (Co-Chair), University of Colombo School of Computing
- Ms. Vidara Ranujee, University of Colombo School of Computing

Registration Committee

- Mr. Viraj Welgama, (Chair), University of Colombo School of Computing
- Ms. Anjalee Devindi, University of Colombo School of Computing
- Ms. Lakshani Gayanthika, University of Colombo School of Computing
- Mr. Tharaka Priyantha, University of Colombo School of Computing
- Mr. S.Thulasigaran, University of Colombo School of Computing
- Ms. Nethmi Keragala, University of Colombo School of Computing
- Ms. Iresha Welikanna,University of Colombo School of Computing

Sponsorship Committee

- Dr. Damitha Sandaruwan, University of Colombo School of Computing
- Dr. Randil Pushpananda, University of Colombo School of Computing

Workshop Handling Committee

- Ms. Sanjani Gunathilaka (Chair), University of Colombo School of Computing
- Ms. Lakshani Gayanthika, University of Colombo School of Computing
- Ms. Sithara Fernando, University of Colombo School of Computing
- Mr. Prasad Wilagama, University of Colombo School of Computing
- Ms. Shashini Wijewardhana, University of Colombo School of Computing
- Ms. Upeksha Sandamini, University of Colombo School of Computing
- Mr. Sumudu Sanjaya , University of Colombo School of Computing

Media and Publicity Committee

- Dr. Samantha Mathara Arachchi, University of Colombo School of Computing
- Mr. Ashintha Perera, University of Colombo School of Computing
- Mr. Shiwantha Wanigathunga, University of Colombo School of Computing
- Mr. Vishad Sudeepa, University of Colombo School of Computing
- Mr. Tharindra Galahena, University of Colombo School of Computing

Logistics Committee

- Mr. Chanaka Colambage, University of Colombo School of Computing
- Mr. F. R. M. Mahroof, University of Colombo School of Computing

Webmaster

- Mr. G. M. T. C. Galahena, University of Colombo School of Computing
- Mr. K. A. R. Perera, University of Colombo School of Computing

Reviewers

- Prof. Amitha Caldera, University of Colombo School of Computing
- Prof. Anuradha Mahasinghe, University of Colombo
- Prof. Anusha Indika, University of Ruhuna
- Prof. Arosha Bandara, Open University of Sri Lanka
- Prof. D.D. Karunaratna, University of Colombo School of Computing
- Prof. David Parry, Murdoch University
- Prof. Hemali Ratnayake, Open University of Sri Lanka
- Prof. Juliana Bowles, University of St Andrews, United Kingdom
- Prof. K.P. Hewagamage, University of Colombo School of Computing
- Prof. Kapila Dias, University of Colombo School of Computing
- Prof. Momina Moetesum, National University of Sciences and Technology
- Prof. Upul J. Sonnadara, University of Colombo
- Dr. C.I. Keppitiyagama, University of Colombo School of Computing
- Dr. C.R. Wijesinghe, University of Colombo School of Computing
- Dr. Duneesha Fernando, University of Melbourne
- Dr. Deepika Seneviratne, University of Sri Jayewardenepura
- Dr. Dharshana Kasthurirathna, University of Colombo School of Computing
- Dr. Dilrukshi Gamage, University of Colombo School of Computing
- Dr. Dilshani Mallikarachchi, Monash University, Malaysia
- Dr. Dinuni Fernando, University of Colombo School of Computing
- Dr. Dinesh Arunatileka, Informatics Institute of Technology
- Dr. Dinusha Vatsalan, SLIIT Northern Uni
- Dr. Dulaji Hidellaarachchi, RMIT University
- Dr. E. Hettiarachchi, University of Colombo School of Computing
- Dr. Fui Chin Hiew, Xiamen University, Malaysia
- Dr. H.N.D. Thilini, University of Colombo School of Computing

- Dr. Hasika Wickrama Senevirathne, Singapore Eye Research Institute, Singapore National Eye Centre
- Dr. Hiran Ekanayake, University of Colombo School of Computing
- Dr. Hiruni Kegalle, University of Melbourne
- Dr. Isuru Dharmadasa, University of the Sunshine Coast
- Dr. Jayalath Ekanayake, Uva Wellassa University
- Dr. Junius Anjana Vidanaralage, SLIIT,
- Dr. K.A.S.H. Kulathilake, Rajarata University of Sri Lanka
- Dr. K.D. Sandaruwan, University of Colombo School of Computing
- Dr. Kasun Gunawardana, University of Colombo School of Computing
- Dr. Kasun Karunanayaka, University of Colombo School of Computing
- Dr. Kaushalya Disanayaka, Brennan
- Dr. Kulani Sumanasekera, New Zealand Skills and Education College
- Dr. Kutila Gunasekera, University of Moratuwa
- Dr. Lasanthi Silva, University of Colombo School of Computing
- Dr. Lojenaa Navanesan, University of Vavuniya
- Dr. M.D.R.N.Dayaratne, University of Colombo School of Computing
- Dr. M.V.P.T. Lakshika, University of Colombo School of Computing
- Dr. Monika Madhavi, University of Colombo
- Dr. Mohammed Moussa,Monash University, Malaysia
- Dr. Nalin Ranasinghe, University of Colombo School of Computing
- Dr. Pradeep Kalansooriya, General Sir John Kotelawala Defence University
- Dr. Randil Pushpananda, University of Colombo School of Computing
- Dr. Roshan Rajapakse, University of Colombo School of Computing
- Dr. S.S.P. Mathara Arachchi, University of Colombo School of Computing
- Dr. Sachintha Pitigala, Whitman College, United States
- Dr. Sanjeewanie Senanayake, Uva Wellassa University
- Dr. Shiromi Arunatileka, University of Colombo School of Computing
- Dr. Subhash Ariyadasa, Uva Wellassa University
- Dr. T. Sritharan, University of Colombo School of Computing
- Dr. T.A. Weerasinghe, University of Colombo School of Computing
- Dr. Thisaranie Kaluarachchi, The Sri Palee Campus of the University of Colombo
- Dr. Thusithanjana Thilakarathna, Auckland University of Technology
- Mr. Akalanka Hettige, Rajarata University of Sri Lanka
- Ms. A.L. Nanayakkara, University of Colombo School of Computing
- Ms. Amali Perera, University of Colombo School of Computing
- Mr. Amod Pathirana, University of Colombo School of Computing
- Mr. Buddhi Gamage, University of Sri Jayewardenepura,
- Mr. Chan Manukalpa, SLIIT City Uni
- Mr. Dasun Bamunuarachchi, University of Colombo School of Computing
- Ms. Dilanga Harshani, University of Colombo School of Computing
- Mr. Eishan Weerasinghe, SLIIT
- Ms. Gayani Wickramarathna, University of Moratuwa

- Mr. Gihan Seneviratne, University of Colombo School of Computing
- Ms. Hansinie Jayathilake, University of Cincinnati, United States
- Ms. Ishani Fonseka, University of Colombo School of Computing
- Mr. Isuru Nanayakkara, University of Colombo School of Computing
- Ms. J.G.S. Rupasinghe, University of Colombo School of Computing
- Ms. Jayathma Chathurangani, University of Colombo School of Computing
- Mr. K.P.M.K. Silva, University of Colombo School of Computing
- Mr. Kenneth Thilakarathna, University of Colombo School of Computing
- Ms. Kalpana Galappaththi, Institute of Technology University of Moratuwa
- Mr. Krishmal Fernando, University of Colombo
- Ms. L.A.S.M. Gunathilaka, University of Colombo School of Computing
- Ms. Lakshani Gayanthika, University of Colombo School of Computing
- Ms. M.A.S.U. Fernando, University of Colombo School of Computing
- Ms. Mathangi Krishnathasan, University of Colombo School of Computing
- Mr. Mohamed Amjath, ATU Donegal, Ireland
- Ms. Nadeesha Epa, University of Colombo School of Computing
- Ms. Nethma Jayawardhane, Asian Institute of Business and Science
- Ms. Nethmi Keragala, University of Colombo School of Computing
- Ms. Nilakshi Shashikala Anuradhani, University of Colombo
- Ms. Nuha Fareed, University of Ruhuna
- Mr. Pasindu Marasinghe, University of Colombo School of Computing
- Mr. Prasad Wilagama, University of Colombo School of Computing
- Mr. Pratheek Senevirathne, University of Melbourne, Australia
- Mr. R.J. Amaraweera, University of Colombo School of Computing
- Mr. R.M.U.A. Rathnayake, University of Colombo School of Computing
- Ms. Rathnaikage Shakya, University of Ruhuna
- Mr. Ravi Madanayake, University of Colombo School of Computing,
- Mr. Roshan Abeyweera, RMIT University, Australia
- Ms. S. Hamsavasini, University of Colombo School of Computing
- Ms. Sanduni Thrimahavithana, University of Colombo School of Computing
- Ms. Sanjeevani Thilakarathne, University of Colombo School of Computing
- Ms. Shashika Ruwanmini, Institute of Technology, University of Moratuwa
- Ms. Thilini Bakmeedeniya, The National School of Business Management
- Mr. Thulasigaran Srimurugathas, University of Colombo School of Computing
- Mr. Veerapathirapillai Vinoharan, University of Vavuniya, Sri Lanka
- Ms. Virajini Godapitiya, Universal College Lanka, Sri Lanka

ORGANIZER

TECHNICAL SPONSORS

SPONSORS

CONFERENCE DAY 1 – AGENDA **Wednesday, 04th November 2026** **Location: Cinnamon Life at City of Dreams, Colombo, Sri Lanka**	
08:00 AM	Registration
08:20 AM	Inauguration and Introduction of ICTer 2026 (About History Video)
08:40 AM	**Welcome Address by the Conference Chair**
08:45 AM	**Address by the UCSC Director**
08:50 AM	**Address by the Chief Guest**
09:00 AM	Cultural Event
09:10 AM	**Panel Discussion: Artificial Intelligence (AI) and the Socio-Economic Future: Challenges and Opportunities**
10:30 AM	Tea Break
11:00 AM	**Oral Presentations Session 1 : Healthcare & Human-Centred Computing**
11:00 AM	NeuLogic-PCOS: A Neuro-Symbolic AI Approach for Explainable PCOS Detection from 2D Ovarian Ultrasound Images
11:15 AM	Evaluating NAO Gestures for Child-Robot Interaction: A Formative Framework Toward Speech Delay Screening.
11:30 AM	NeuroAgent: A Brain-to-Device Communication Framework Integrating c-VEP EEG Decoding with Agentic AI for Desktop Task Execution
11:45 AM	Cloud-Based Distributed Data Processing: Scalable Data-partitioning Techniques for Healthcare-data
12:00 PM	Certification Ceremony for Oral Presentations Session 1
12:10 PM	Title Sponsor Message: Mastercard
12:30 PM	Lunch Break
01:30 PM	**Keynote - Dr. René E. D. Ferdinands** Topic: Engineering the Future of Spin Bowling: Biomechanics, Smart Ball Analytics and 21st-Century Coaching

02:15 PM	**Poster Presentations Session 1: Computer Vision & Graphics**
02:15 PM	Text2NeRF: Multi-Aspect Knowledge Distillation in Diffusion-Guided NeRFs for Text-Driven 3D Synthesis
02:20 PM	Deployment-Aware Optimization of Real-Time Neural Rendering Under Fixed Runtime Constraints
02:25 PM	CloneVerify: Quality-Aware Conditional Enhancement for Fine-Grained Sri Lankan Tea Clone Identification
02:30 PM	Small Object Detection Enhancement in Real-Time Detection Transformer (RT-DETR) for UAV Imagery
02:35 PM	Certification Ceremony for Poster Presentations Session 01
02:45 PM	**Poster Presentations Session 2: Intelligent Sensing, IoT & Adaptive Learning**
02:45 PM	Towards Non-Visual Olfactory Sensing: An Intelligent Electronic Nose for Detecting Sri Lankan Elephants (Elephas maximus maximus)
02:50 PM	Development of an Intelligent and Secure Smart Home Automation Using Home Assistant and IoT Devices
02:55 PM	Quantum-Aware Adversarial Defense for IoT Intrusion Detection Systems Using Explainable AI
03:00 PM	SmartFedProx: Personal Finance Modelling Using Federated Learning to Address Data Heterogeneity
03:05 PM	A Hybrid Concept Drift-Aware Streaming Learning for Real-Time IT Project Risk Prediction
03:10 PM	Certification Ceremony for Poster Presentation Session 02
03:15 PM	Tea Break
03:45 PM	Industry Talk - Commercial Bank
04:00 PM	**Industry Track Presentations**
04:00 PM	Quantum Risk Assessment Framework for Harvest Now, Decrypt Later (HNDL) Attacks in Banking Sector
04:15 PM	An AI-Enabled Digital Learning Passport with Skill Gap Detection and Industry-Integrated Recruitment Interface: A Design Science Approach
04:30 PM	An Edge-Cloud Machine Condition Monitoring System for Industrial Rubber Processing

04:45 PM	Certification Ceremony for Industry Track Presentations
5:00 PM	Day one concluding

CONFERENCE DAY 2 – AGENDA Thursday, 5th November 2026 Location: Cinnamon Life at City of Dreams, Colombo, Sri Lanka	
08:00 AM	Registration
08:30 AM	**Oral Presentation Session 2 : Optimization, Learning & Autonomous Systems**
08:30 AM	Forecasting Foreign Exchange Reserves Through Crisis and Recovery: Deep Learning Evidence from Sri Lanka
08:45 AM	Quantum Mutual Information for Causal Feature Selection: An Empirical Study of Feature Recovery and Treatment Effect Estimation
09:00 AM	Impact of Adopting Features of Recent Nature-Inspired Algorithms on the Behaviour of the Simple Genetic Algorithm
09:15 AM	Hybrid Adaptive Coordination in Multi-UAV Swarms for Search and Rescue
09:30 AM	Certification Ceremony for Oral Presentation Session 2
09:45 AM	**Keynote – Mr. Sampath Tilakumara, London Stock Exchange Group (LSEG)** Topic: Who Gets Left Behind? AI, Work, and the Next Generation of Talent
10:30 AM	Tea Break
11:00 AM	**Oral Presentation Session 3 : Cloud, Distributed Systems & Resource Management**
11:00 AM	Application-Aware Dynamic Virtual Machine Placement for Live Migration
11:15 AM	HAQL-PR: A Hybrid Adaptive Q-Learning Page Replacement Algorithm for Efficient Memory Management
11:30 AM	Analyzing Organizational Influence on Microservice Communication Affinity and Designing Adaptive Colocation Strategies
11:45 AM	D-Track: Eliminating Redundant Fake Dirty Pages in Live Virtual Machine Migration
12:00 PM	Certification Ceremony for Oral Presentation Session 3

12:10 PM	Title Sponsor Message: Mastercard
12:30 PM	Lunch Break
01:30 PM	**Keynote – Dr. Buddy Liyanage, University of Colombo School of Computing** Title:The need for responsible, ethical, and trustworthy AI systems and their use in Sri Lanka
02:15 PM	**Poster Presentations Session 3: Applied Data Analytics and Machine Learning**
02:15 PM	SinFair: Mitigating Demographic Entity Erasure in Sinhala Abstractive Summarization via Five-Axis Counterfactual Augmentation
02:20 PM	Federated Reasoning: A Privacy-Preserving Medical Question Answering System with Chain-of-Thought Reasoning and Retrieval-Augmented Generation
02:25 PM	Quantifying Metformin Confounding in Machine Learning-Based Metagenomic Glycemic Classification
02:30 PM	Attitudinal Predictors of Music Genre Exploration and Abandonment Among Streaming Users: A Survey-Based Multi-method Analysis
02:35 PM	Certification Ceremony for Poster Presentations Session 3
02:45 PM	**Poster Presentations Session 4: Software and Systems**
02:45 PM	Enhancing Live VM Migration for Latency-Sensitive Workloads
02:50 PM	When the Next Step Is Not One Step: Distribution-Aware Execution Modeling for Concurrent Go Programs
02:55 PM	Certification Ceremony for Poster Presentations Session 4
03:15 PM	Tea Break
03:45 PM	**Oral Presentation Session 4: LLMs, Agentic AI & Trustworthy AI**
03:45 PM	Beyond Correctness: Output and Prompt Bloat in LLM-Generated Code
04:00 PM	Reasoning Path Hijacking in Small Language Models: Vulnerability Assessment for Agentic AI Systems
04:15 PM	Knowledge Graph-Based Fact Inclusion Explanation for Abstractive Summarization
04:30 PM	Certification Ceremony for Oral Presentation Session 4
04:35 PM	Awards Ceremony

04:40 PM	Finale and Next Year Promo Video
04:50 PM	Vote of Thanks
05:00 PM	Day two concluding

KEYNOTES

INTRODUCTION TO THE KEYNOTE SPEAKER 1



Dr. René E. D. Ferdinands

Cricket Biomechanics Specialist
International Cricket Coach
Sri Lanka / Australia / New Zealand / India

Dr. René E. D. Ferdinands is an international cricket coach, former first-class cricketer, and internationally recognised cricket biomechanics specialist. He recently served as National Spin Bowling Coach for Sri Lanka Men's Cricket and has worked across elite cricket environments in Sri Lanka, Australia, New Zealand, and India. He brings extensive expertise in cricket biomechanics, spin-bowling mechanics, injury-risk analysis, bowling-action assessment, smart cricket ball technology, and performance analytics. He has published extensively in peer-reviewed sport science and sports engineering literature and has taught elite players, coaches, and coach educators internationally.

ABSTRACT OF THE KEYNOTE ADDRESS 1

Engineering the Future of Spin Bowling: Biomechanics, Smart Ball Analytics and 21st-Century Coaching

Dr. René E. D. Ferdinands

University of Sydney

Spin bowling remains one of cricket's most complex skill systems, integrating whole-body biomechanics, finger-wrist torque production, rotational physics, tactical deception and long-term motor learning. This keynote presents a 21st-century scientific framework for understanding and developing spin bowling by integrating motion analysis, smart cricket ball technology and coaching intelligence. It examines the kinematic and dynamic signatures that differentiate variant forms of spin delivery, including their effects on spin rate, spin-axis orientation, precession, torque generation, efficiency and ball behaviour.

A key component is the advanced smart cricket ball, an instrumented ball equipped with high-speed gyroscopic sensing that captures rotational performance variables not readily observable through video alone. When smart ball telemetry is combined with motion analysis and biomechanical modelling, coaches and researchers can move beyond visual technique correction toward profiling of delivery types, skill efficiency, performance limitations and training effects. This creates a practical bridge between laboratory biomechanics and field-based coaching.

The presentation discusses applications in coaching interventions, action remediation, talent identification, delivery repertoire development and the monitoring of training outcomes. It also positions spin bowling as a rich cyber-physical problem in which human movement, embedded sensing, data analytics, physics-based modelling and decision support converge in cricket.

Future directions include smart coaching systems, data-driven delivery classification, group-theoretic models of spin variation and advanced visual demonstrations of elite spin mechanics. Ultimately, the keynote argues that the future of spin bowling lies not in replacing artistry, but in making its hidden mechanics measurable, coachable and transferable to match performance.

INTRODUCTION TO THE KEYNOTE SPEAKER 2



Mr. Sampath Tilakumara

Director,
Software Engineering and Distinguished Engineer,
London Stock Exchange Group (LSEG)

Sampath Tilakumara is “Director, Software Engineering” and a Distinguished Engineer at the London Stock Exchange Group (LSEG), with over 22 years of experience delivering ultra-low-latency trading systems and mission-critical capital market technologies. He currently works on the evaluation, development, and rollout of Agentic AI frameworks and AI-driven Software Development Lifecycle (AI-SDLC) practices at LSEG, bringing a practitioner's perspective on adopting AI responsibly in high-stakes, regulated environments.

ABSTRACT OF THE KEYNOTE ADDRESS 2

Who Gets Left Behind? AI, Work, and the Next Generation of Talent

Mr. Sampath Tilakumara

London Stock Exchange Group (LSEG)

We keep asking whether AI will take our jobs. The more urgent question is who is being prepared to work alongside it - and who isn't.

This keynote argues that three shifts are colliding faster than our institutions can adapt. First, the shape of work itself is changing: as AI systems move from doing tasks for us to working with us autonomously, routine execution loses value while judgement, specification, verification, and orchestration gain it. Software is the leading edge of this pattern, not the whole of it.

Second, that transition is not evenly available. The real divide is not access to AI tools but access to the expertise to use them well - leaving smaller firms and developing economies at risk of becoming permanent consumers, rather than producers, of AI-driven value.

Third, and most consequentially, our education and hiring systems are the pipeline that decides who arrives ready. Curricula still train and test for skills AI now performs trivially. Assessment will have to shift from "can you produce the answer" to "can you frame the problem, judge the output, and catch the failure" - and employers will follow, valuing evaluative literacy over output volume.

For academics, this is an assessment design problem. For industry, a hiring problem. For policymakers, it is where national competitiveness will be won or lost.

INTRODUCTION TO THE KEYNOTE SPEAKER 3



Dr. Buddy Liyanage

Senior Visiting Lecturer,
University of Colombo School of Computing (UCSC),
Responsible, Ethical, and Trustworthy AI Governance Advisor,
Independent AI Researcher

Dr. Buddy Liyanage has been serving as a Senior Visiting Lecturer and Advisor to the University of Colombo School of Computing (UCSC) for over twenty years. He brings extensive expertise in cybernetics, networking, cybersecurity, and security architecture. He received his BSc in Physics (Honours) from the University of Colombo, followed by an MSc in Modern Electronics from the University of Nottingham. He subsequently completed his PhD in Cybernetics at the University of Reading in the United Kingdom. He has also served as the Head of Security Architecture for over seventeen years.

ABSTRACT OF THE KEYNOTE ADDRESS 3

The need for responsible, ethical, and trustworthy AI systems and their use in Sri Lanka

Dr. Buddy Liyanage

University of Colombo School of Computing (UCSC)

Sri Lanka stands at a significant point in its national digital transformation, driven by policy ambitions to expand its digital economy to US$15 billion and establish the island as a regional technology hub by 2030. Rapid artificial intelligence (AI) use presents dual realities: promising innovations (such as AI-assisted diabetic retinopathy screening in low-resource regional clinics) alongside significant societal vulnerabilities (such as political deep fakes, investment fraud, and automated misinformation). This keynote address argues that responsibility, ethics, and trustworthiness are not optional additions to digital development but essential engineering requirements without which technology alone will not deliver sustainable value.

The address establishes formal conceptual precision by delineating ethical AI (normative human values), responsible AI (institutional accountability and governance), and trustworthy AI (verifiable technical reliability), while addressing the global principles-to-practice gap in a Sri Lankan context. Using grounded Sri Lankan case studies spanning healthcare diagnostics, social welfare eligibility, agricultural crop advice, child protection in education, sustainable higher education, democratic integrity, and local-language model performance, the speech demonstrates how algorithmic bias, opacity, and linguistic disparities in Sinhala and Tamil can impact society. It highlights the foundational legislative anchors in Sri Lanka's Personal Data Protection Act (PDPA No. 9 of 2022), which grant citizens explicit statutory rights regarding automated decision-making.

To translate policy vision into operational readiness, the address concludes by outlining a practical national action roadmap with multiple points. Key recommendations include integrating mandatory AI risk evaluations into Data Protection Impact Assessments (DPIAs), enforcing responsible public procurement standards for AI, establishing university-hosted Sinhala and Tamil evaluation test beds, guaranteeing citizens' statutory rights to explanation and administrative appeal, and fostering civic AI literacy. Ultimately, the speech challenges researchers, policymakers, and industry leaders to evaluate national progress not merely by the volume of deployed algorithms and AI systems, but by the depth of public trust in such systems, demonstrating AI trustworthiness in Sri Lanka.

Keywords: *Responsible AI, Ethical AI, Trustworthy AI, AI Governance, Sri Lanka*

TABLE OF CONTENTS

CONFERENCE PAPERS

- **TECHNICAL SESSION 1**

NeuLogic-PCOS: A Neuro-Symbolic AI Approach for Explainable PCOS Detection from 2D Ovarian Ultrasound Images
Amaya De Silva[1,2] ,Janith Prabhanuka[1]* 30

Evaluating NAO Gestures for Child-Robot Interaction: A Formative Framework Toward Speech Delay Screening
Chan Sri Manukalpa, Koliya Pulasinghe, and Samantha Rajapaksha* 31

NeuroAgent: A Brain-to-Device Communication Framework Integrating c-VEP EEG Decoding with Agentic AI for Desktop Task Execution
Ashen Himeshana[1], Janith Prabhanuka[1]* 32

Cloud-based Distributed Data processing: Scalable Data-partitioning Techniques for Healthcare-data
Vishmi Wickramarachchi[1], Sanduni Thilakarathne[1], and Jananie Jarachanthan[1]* 33

- **TECHNICAL SESSION 2**

Forecasting Foreign Exchange Reserves Through Crisis and Recovery: Deep Learning Evidence from Sri Lanka
Yasiru Rasintha Dassanayake[1], Samantha Mathara Arachchi[1], Ishani Fonseka[1], Hiruni Navodya Kudagama[1]* 35

Quantum Mutual Information for Causal Feature Selection: An Empirical Study of Feature Recovery and Treatment Effect Estimation
Janith R.R.H. Ramanayake[1]* 36

Impact of Adopting Features of Recent Nature-Inspired Algorithms on the Behaviour of the Simple Genetic Algorithm
Bellanthuda Achchige Senali Dilumika[1],,Gihan Pushpajith Seneviratne[1]* 37

Hybrid Adaptive Coordination in Multi-UAV Swarms for Search and Rescue
Pramuditha Vithanaarachchi[1,2], Janith Prabhanuka[1], and Thathsara Nanayakkara[1,3]* 38

- **TECHNICAL SESSION 3**

Application-Aware Dynamic Virtual Machine Placement for Live Migration
Senuri Wickramasinghe[1], Pratheek Senevirathne[1], Jerome Dinal Herath[1], Dinuni K. Fernando[1]* 40

HAQL-PR: A Hybrid Adaptive Q-Learning Page Replacement Algorithm for Efficient Memory Management
Sangeetha Arunprakash[1], Vishaliney Pirathap[1], Sinthujan Ragunathan[1], Samantha Rajapaksha[2]* 41

Analyzing Organizational Influence on Microservice Communication Affinity and Designing Adaptive Colocation Strategies
Sella Hewage Vinupa Thenuka[1], Kande Ranasinghage Kasun Hansamal[1], Liyanage Danujaya Sandeepa Liyanage[1], Chamath Keppitiyagama[1], and Tharindu Wijethilake[1]* 42

D-Track: Eliminating Redundant Fake Dirty Pages in Live Virtual Machine Migration
Samindu R. Cooray[†] , T. Kasun O. Mendis[†], Dinuni K. Fernando[*].....43

- **TECHNICAL SESSION 4**

Beyond Correctness: Output and Prompt Bloat in LLM-Generated Code
Pasandi Ranganath[1], Chamath Keppitiyagama[1],Tharindu Wijethilake[1]*.....45
Reasoning Path Hijacking in Small Language Models: Vulnerability Assessment for Agentic AI Systems
Sandares Fernando[1] , Randil Pushpananda[1], Ruvan Weerasinghe[2]*..... 46
Knowledge Graph-Based Fact Inclusion Explanation for Abstractive Summarization
Nirasha Nelki[*], Dilrukshi Gamage.....47

- **POSTER SESSION 1**

Text2NeRF: Multi-Aspect Knowledge Distillation in Diffusion-Guided NeRFs for Text-Driven 3D Synthesis
Thisarani Wickramasinghe[1], RGN Meegama[2].....49
Deployment-Aware Optimization of Real-Time Neural Rendering Under Fixed Runtime Constraints
B.H.S. Tharuka Fernando[1], P.P.G. Dinesh Asanka[1]*.....50
CloneVerify: Quality-Aware Conditional Enhancement for Fine-Grained Sri Lankan Tea Clone Identification
Sadesh Weerasuriya[1], Imesh Pathirana[1]*..... 51
Small Object Detection Enhancement in Real-Time Detection Transformer (RT-DETR) for UAV Imagery
J.M. Sasanka Madhura Bandara [(*)] , Pumudu A. Fernando..... 52

- **POSTER SESSION 2**

Towards Non-Visual Olfactory Sensing: An Intelligent Electronic Nose for Detecting Sri Lankan Elephants (Elephas maximus maximus)
W.M.A.T Wanninayake[1] , P.W Paranawithanage[1] , T.S Dissanayake[1], W.M.H.K Weerakoon[1], R.G.B.C Rajagalgoda[1], B.M.G.S.T.S.K Wickramasinghe[1] [*], N.A.A.M Prabhath[1], G.D Illeperuma[1], T.S.P Fernando[1].....54
Quantum-Aware Adversarial Defense for IoT Intrusion Detection Systems Using Explainable AI
Supun P. Shaminda[1,], Charaka Geethal Kapugama[1]*.....55
SmartFedProx: Personal Finance Modelling Using Federated Learning to Address Data Heterogeneity
V A D Jeremy Perera[1], Torin Wirasingha[2]*..... 56
A Hybrid Concept Drift-Aware Streaming Learning for Real-Time IT Project Risk Prediction
W. K. C. P. Bandara[*], A. M. C. Malkanthi.....57

- **POSTER SESSION 3**

SinFair: Mitigating Demographic Entity Erasure in Sinhala Abstractive Summarization via Five-Axis Counterfactual Augmentation
Anupa Wanigasinghe[1], Salitha Perera[1]* 59
Federated Reasoning: A Privacy-Preserving Medical Question Answering System with Chain-of-Thought Reasoning and Retrieval-Augmented Generation
J. G. Sevindu Imsara[1,2], Janith Prabhanuka[1]* 60
Quantifying Metformin Confounding in Machine Learning-Based Metagenomic Glycemic Classification
Ruchila Wijerathne[1], Rupika Wijesinghe[1]* 61
Attitudinal Predictors of Music Genre Exploration and Abandonment Among Streaming Users: A Survey-Based Multi-method Analysis
Weerasuriya Arachchige Malsha Prabuddhi[1], Saminda Chandika Premaratne[2] , Podduwa Kankanamage Subash Chaminda Jayasinghe[1]* 62

- **POSTER SESSION 4**

Enhancing Live VM Migration for Latency-Sensitive Workloads
Ayash Maddumage[1] , Nadeesha Epa[1] , Dinuni K. Fernando[1]* 64
When the Next Step Is Not One Step: Distribution-Aware Execution Modeling for Concurrent Go Programs
Kaviru Hapuarachchi[1] 65

- **INDUSTRY TRACK**

Development of an Intelligent and Secure Smart Home Automation Using Home Assistant and IoT Devices
[1]Rumesh Rajasegar, [2]Thayaparan Subramaniam, [2]Mukunthan Tharmakulasingam, ...67
Quantum Risk Assessment Framework for Harvest Now, Decrypt Later (HNDL) Attacks in Banking Sector
*Vishvadini Kurukulasooriya[*1], S. Kasun Chathuranga[1], Udara Wijesinghe[2]* 71
An AI-Enabled Digital Learning Passport with Skill Gap Detection and Industry-Integrated Recruitment Interface: A Design Science Approach
Madusanka Premaratne Rathnayake Mudiyanselage[1] , Hasanthi Lakmali Thellapura Arachchilage[1] , Dimuthma Umashani1 ,Dillina Herath[2]* 72
An Edge-Cloud Machine Condition Monitoring System for Industrial Rubber Processing
Rumesh Rajasegar 73

TECHNICAL SESSION 1

Healthcare & Human-Centred Computing

NeuLogic-PCOS: A Neuro-Symbolic AI Approach for Explainable PCOS Detection from 2D Ovarian Ultrasound Images

Amaya De Silva[1,2] ,Janith Prabhanuka[1*]

[1] *School of Computing, Informatics Institute of Technology,Sri Lanka,*

[2] *University of Westminster, London, United Kingdom*

Polycystic Ovary Syndrome (PCOS) is a prevalent endocrine disorder affecting 6-13% of women of reproductive age globally, with up to 70% remaining undiagnosed. While deep learning models for PCOS detection from ovarian ultrasound images have achieved high accuracy, they operate as black-box systems that cannot provide clinically grounded reasoning, limiting adoption in clinical settings. This paper presents NeuLogic-PCOS, a neuro-symbolic AI framework that integrates a fine-tuned EfficientNetB0 Convolutional Neural Network (CNN) with a CART Decision Tree symbolic reasoning layer built on Rotterdam-aligned morphological features in the Rotterdam diagnostic criteria. Six morphological features selected using Cohen's d effect size (all $d > 1.0$) are extracted via Laplacian of Gaussian (LoG) blob detection and represent image-derived indicators of polycystic ovarian morphology (PCOM). A weighted ensemble ($\alpha = 0.5$, threshold $\theta = 0.6936$), jointly optimised on a deduplicated validation set, combines neural image perception with rule-based clinical reasoning: LIME-based visual explanations and a three-level confidence assessment module further support clinical transparency. Evaluated on a deduplicated combined dataset of approximately 4,000 images, on a held-out test set of 597 images, NeuLogic-PCOS achieves 97.65% accuracy, 95.98% recall, 98.85% specificity, 97.15% F1-score, and AUC 0.9870, outperforming standalone CNN (96.48%) and symbolic-only (96.98%) baselines. NeuLogic-PCOS is intended as an ultrasound-based PCOM screening aid and does not replace the clinical assessment required for PCOS diagnosis under the Rotterdam criteria.

Keywords: *PCOS detection, Neuro-Symbolic AI, Ovarian Ultrasound*

Evaluating NAO Gestures for Child-Robot Interaction: A Formative Framework Toward Speech Delay Screening

Chan Sri Manukalpa*, Koliya Pulasinghe, and Samantha Rajapaksha
Department of Information Technology, Sri Lanka Institute of Information Technology, Malabe, Sri Lanka

Socially assistive robots offer a setting for culturally adapted interaction, but sustained robotic choreography remains underexplored in Sri Lanka. This study presents a framework for designing and formative evaluating NAO gestures, with speech-delay screening as a planned future application. Two routines represented contrasting design cases: Kandyan-inspired cultural familiarity and Hip-hop-inspired rhythmic dynamism. Development combined Webots simulation, physical execution, and 12 preliminary interaction trials with typically developing children, divided equally between the routines. These trials supported troubleshooting and informal feasibility assessment; attention and speech outcomes were not measured. A six-member professional panel appraised the routines, while separate parent (N = 20) and undergraduate (N = 228) groups rated silent recordings. Kandyan received higher expert-rated cultural appropriateness (M = 4.72, SE = 0.15) and parent-rated perceived engagement (M =4.42, SE = 0.20), whereas Hip-hop received higher undergraduate-rated dynamism (M = 4.46, SD = 0.28). Expert appraisal informed refinement; adult ratings described perceived qualities of the recordings. The findings inform further development of culturally grounded NAO choreography. Children with Autism Spectrum Disorder (ASD) were not included, and screening performance, ASD-related attentional benefit, and therapeutic effectiveness remain untested. Adaptive selection and automated library acceptance are presented separately as future specifications.

Keywords: *Autism Spectrum Disorder, Socially Assistive Robotics, Human Robot Interaction, Cultural Adaptation, Robotic Choreography, NAO Robot, Assistive Technology*

NeuroAgent: A Brain-to-Device Communication Framework Integrating c-VEP EEG Decoding with Agentic AI for Desktop Task Execution

Ashen Himeshana[1*], Janith Prabhanuka[1]

[1]*Department of Computing, Informatics Institute of Technology, Sri Lanka*

Most EEG-based brain-computer interfaces (BCIs) map neural signals to a small, fixed set of commands, whereas large language model (LLM) agents can operate real graphical interfaces but receive typed instructions. NeuroAgent connects the two. Selections decoded from code-modulated visual evoked potentials (c-VEP) with canonical correlation analysis (CCA) drive a perception-based agent that parses the screen with OmniParser, uses a locally hosted LLM (Gemma3:4b) to turn detected elements into goal-labelled options, executes the chosen option, and verifies the outcome, using Win32 window titles for application launches and a vision-language check for in-app actions, before replanning. We report two evaluations. First, the decoder was evaluated offline on the real 32-channel EEG of the public VP1 c-VEP dataset (96 held-out trials, 960 six-way decisions). With filter and lag parameters selected on the training split only, accuracy was 97.4% (95% CI 96.2–98.2%) with 3 s windows, whereas the prototype's original reference construction performed at chance (15.9%). Second, an integration demonstration in which decoded selections drove the agent on an unmodified Windows desktop completed 14 of 15 goals across four application contexts; this demonstration used synthetic EEG with simulator-chosen targets and is not a measure of intent-driven performance. We also compare four small local LLMs for the option-generation step. The results form a single-participant, offline proof of concept; a live closed-loop user study is required before claims about practical use can be made.

Keywords: *Brain-computer interface · c-VEP · EEG decoding · Agentic AI · GUI agents · Assistive technology*

Cloud-based Distributed Data processing: Scalable Data-partitioning Techniques for Healthcare-data

Vishmi Wickramarachchi[1*], Sanduni Thilakarathne[1], and Jananie Jarachanthan[1]

Faculty of Engineering, University of Jaffna, Sri Lanka

Cloud computing has become a core technology for health care, offering virtually unlimited storage and elastic compute to manage rapidly growing clinical datasets such as electronic health records, medical images, and IoT sensor streams. However, traditional monolithic processing models struggle with these workloads, leading to performance bottlenecks, load imbalance, and scalability limits. This research investigates scalable data-partitioning techniques for cloud-based distributed data processing, with a specific focus on healthcare data. Building on recent advances in distributed frameworks, sharding, and parallel analytics, the work evaluates multiple partitioning strategies and their impact on scalability, load distribution, latency, and overall system efficiency in healthcare scenarios. Here, the raw datasets collected from wearable devices are cleaned, formatted, and normalized. The study considers heterogeneous data types and access patterns common in medical environments, such as patient-centric queries, real-time monitoring, and AI-driven diagnostics. A cloud-native experimental framework is designed to split large healthcare datasets into manageable partitions and process them in parallel across distributed resources. Quantitative evaluation demonstrates how appropriate partitioning improves throughput, reduces response time, and better utilizes compute and storage resources compared with non-partitioned or naively partitioned baselines.

Keywords: *Distributed data processing, hash partitioning, range partitioning, directory partitioning, hybrid partitioning.*

TECHNICAL SESSION 2

Optimization, Learning & Autonomous Systems

Forecasting Foreign Exchange Reserves Through Crisis and Recovery: Deep Learning Evidence from Sri Lanka

Yasiru Rasintha Dassanayake[1*], Samantha Mathara Arachchi[1], Ishani Fonseka[1], Hiruni Navodya Kudagama[1]

University of Colombo School of Computing, Sri Lanka

Foreign-exchange reserves are a central indicator of external stability and a core input to monetary and reserve-management policy. Forecasting them is difficult when reserves move from accumulation through crisis depletion into rebuilding, as Sri Lanka did over 2019-2025. To our knowledge, this is the first pseudo-real-time comparison of accounting, econometric, regime-switching, machine-learning, and deep-learning forecasts of a reserve series through a full accumulation-default-rebuilding cycle, evaluated against both a statistical and an economic (balance-of-payments identity) benchmark, with errors expressed in months of import cover. All models are evaluated in a pseudo-real time, expanding-window rolling-origin design in which each forecast uses only information available at the origin. The balance-of-payments accounting identity is the strongest one-month forecaster, showing that near-term reserve movements are best captured by accounting mechanics and level persistence rather than more complex learned structure. At three and six months, the LSTM leads the model set, delivering meaningful skill against both the random walk and the identity benchmark, with forecast errors remaining below one month of import cover; its accuracy deteriorates when its input window reaches back across the 2022 crisis. At longer horizons, forecast accuracy deteriorates enough to limit operational use. Regime-switching and tree-based models do not deliver a general advantage in this window. Reserve monitoring should therefore pair the accounting identity for next-month tracking with recurrent models for three-to-six-month planning.

Keywords: *Foreign exchange reserves · Forecasting · LSTM · Structural breaks · Sri Lanka*

Quantum Mutual Information for Causal Feature Selection: An Empirical Study of Feature Recovery and Treatment Effect Estimation

Janith R.R.H. Ramanayake[1*]

[1]*Department of Statistics, University of Colombo, Sri Lanka*

We investigate whether Quantum Mutual Information (QMI), defined through the von Neumann entropy of angle-encoded density matrices, provides a useful lens on causal feature selection. Using six synthetic data-generating processes (DGPs) and an IHDP inspired synthetic benchmark, we compare QMI-based confounder selection with information-theoretic, kernel (HSIC), regularization and tree-based selectors, and ask how well feature recovery predicts downstream average treatment effect (ATE) estimation. Our central finding is that confounder-recovery F1 is only a weak predictor of ATE error: it explains about 4% of the variance across DGPs and 8–13% on identical simulation draws, and within draws the association vanishes once $n \geq 500$. On the IHDP-inspired benchmark, methods with F1 between 0.51 and 1.00, including QMI-CFS, are statistically tied in ATE error. We argue that this weak coupling is expected from adjustment-set theory and recommend evaluating causal feature selectors jointly on recovery and estimation. As a selector, QMI is competitive with k-NN mutual information but is outperformed in recovery by tree-based wrappers on four of six DGPs. We find no evidence of quantum advantage: QMI is best understood as a bandwidth-free, fixed-kernel dependency estimator.

Keywords: *Quantum Mutual Information · Causal Inference · Feature Selection · Treatment Effect Estimation · Quantum-Inspired Methods*

Impact of Adopting Features of Recent Nature-Inspired Algorithms on the Behaviour of the Simple Genetic Algorithm

Bellanthuda Achchige Senali Dilumika[1,*],Gihan Pushpajith Seneviratne[1]

[1]*University of Colombo School of Computing, Sri Lanka*

Hybrid extensions of the Simple Genetic Algorithm (SGA) are usually judged only by final error, leaving it unclear how a grafted feature changes the underlying search. We hold the SGA fixed and swap one operator at a time, producing sixteen single graft variants: twelve from six nature inspired families and four operator research controls. For each variant we measure how the graft changes the search itself: how widely the population explores, how much diversity it keeps and how its mutations behave, alongside final error on CEC 2022 at D ∈ {10, 20}, five CEC 2020 engineering problems and seven TSPLIB instances (51–442 cities) with paired Wilcoxon, Cliff's δ and Friedman tests. The grafts split into clear winners and losers and grafts from the same family fall on both sides (SGA-SHADE and SGA-DE), so the source metaphor is a poor guide to a graft's effect. Runaway/divergent grafts degrade CEC 2022 by a large margin at both D = 10 and D = 20. Among the D = 10 winners, pheromone memory (SGA-Pher) and current-to-pbest mutation (SGA-SHADE) lead on large TSP and stay clear of the divergent grafts at D = 20, while the diversity injection wins (SGA-NetSup, SGA-Cuckoo) hold only inside the smaller envelope. TSP results for natively continuous operators pass through an edge-aware mapping and should be read accordingly. We describe what we observe and do not claim a general taxonomy.

Keywords: *Simple Genetic Algorithm · Nature inspired algorithms · Behavioural analysis · Metaheuristic*

Hybrid Adaptive Coordination in Multi-UAV Swarms for Search and Rescue

Pramuditha Vithanaarachchi[1,2], Janith Prabhanuka[1*], and Thathsara Nanayakkara[1,3]

[1]School of Computing, Informatics Institute of Technology, Sri Lanka, [2] Robert Gordon University, Scotland,UK, [3]Industry 4.0 Convergence Bionics Engineering, Pukyong National University, Republic of Korea

Search and rescue (SAR) operations require rapid exploration of complex and hazardous environments, where efficient coordination directly impacts mission success. While multi-UAV swarms offer advantages in coverage and robustness, coordinating multiple agents to balance exploration, navigation, and response remains a key challenge. This paper presents a phase-based decentralized coordination framework for multi-UAV SAR systems, integrating swarm intelligence-based exploration, obstacle-aware navigation, and adaptive role-based task allocation within a unified architecture. The framework enables progressive integration and controlled comparison of coordination strategies with increasing coordination complexity. The proposed system is evaluated in the Webots simulation environment across different terrain sizes and swarm configurations, with performance assessed using mission-oriented metrics including mission duration, coverage, rescue time, communication overhead, and navigation efficiency. Results show that no single coordination strategy consistently outperforms the others across all evaluated metrics; instead, the strategies exhibit distinct trade-offs between efficiency, rescue responsiveness, coordination complexity, and communication requirements. The findings demonstrate the value of controlled within-framework evaluation for examining the effects of progressively integrated coordination mechanisms and provide insights into the trade-offs involved in designing adaptive multi-UAV SAR systems.

Keywords: *Multi-UAV Systems · Swarm Intelligence · Multi-Agent Coordination · Search and Rescue · Decentralized Coordination · Simulation-Based Evaluation*

TECHNICAL SESSION 3

Cloud, Distributed Systems & Resource Management

Application-Aware Dynamic Virtual Machine Placement for Live Migration

Senuri Wickramasinghe[1], Pratheek Senevirathne[1], Jerome Dinal Herath[1], Dinuni K. Fernando[1*]

[1]*University of Colombo School of Computing, Sri Lanka*

Dynamic Virtual Machine Placement (VMP) often ignores the communication dependencies of multi-tier cloud applications by treating VMs as isolated entities. To address this limitation, this work presents the Application-Aware Virtual Machine Placement (AAVMP) framework. Formulated as a Markov Decision Process, AAVMP uses a Graph Neural Network (GNN) to encode real-time inter-VM traffic and physical rack topology. A Mixture of Experts (MoE) Deep Q-Network then routes decisions through specialized pathways to balance host resource consolidation and topological co-location dynamically. In addition, a multi-objective reward function with deferred batch correction enables accurate training during concurrent migrations. The proposed model was evaluated against six baselines, where AAVMP achieved a 100% topology-aware co-placement rate with zero migration failures, the lowest application latency (1.670 s), and the lowest weighted communication cost (1.132). AAVMP shows scalability on infrastructures with up to 320 hosts, consistently co-optimizing performance, safety, and infrastructure efficiency.

Keywords: *virtual machine placement, deep reinforcement learning, graph neural network, mixture of experts, live migration, cloud computing*

HAQL-PR: A Hybrid Adaptive Q-Learning Page Replacement Algorithm for Efficient Memory Management

Sangeetha Arunprakash[1*], Vishaliney Pirathap[1], Sinthujan Ragunathan[1], Samantha Rajapaksha[2]

[1]Department of Computing, SLIIT Northern Uni, Sri Lanka, [2]Faculty of Computing, Sri Lanka Institute of Information Technology, Sri Lanka

Efficient page replacement is important for improving virtual memory performance and reducing page faults in operating systems. Classical algorithms such as FIFO, LRU, and CLOCK are simple and efficient, but they use fixed rules and may not adapt well to changing memory access patterns. This paper proposes HAQL-PR, a Hybrid Adaptive Q-Learning Page Replacement algorithm that combines recency information, access frequency, and Q-learning-based eviction values within a Markov Decision Process framework. A Python-based simulator is used to compare HAQL-PR with FIFO, LRU, CLOCK, and pure Q-learning under four frame configurations: 3, 5, 8, and 10 frames. The results show that HAQL-PR achieves the lowest page fault count and the highest hit ratio, reaching 66.1% at 10 frames. The proposed method also converges faster than pure Q-learning with moderate computational overhead.

Keywords: *Page Replacement, Virtual Memory, Q-Learning, Reinforcement Learning, HAQL-PR*

Analyzing Organizational Influence on Microservice Communication Affinity and Designing Adaptive Colocation Strategies

Sella Hewage Vinupa Thenuka[1*], Kande Ranasinghage Kasun Hansamal[1], Liyanage Danujaya Sandeepa Liyanage[1], Chamath Keppitiyagama[1], and Tharindu Wijethilake[1]

University of Colombo School of Computing, Sri Lanka

Tightly coupled microservice interactions, manifesting as Blob and Empty-semi-trucks anti-patterns, impose measurable network overhead that traditional observability tools cannot detect without code instrumentation or sidecar proxies. This paper presents an automated, non-disruptive framework that captures inter-service telemetry at the kernel level using extended Berkeley Packet Filter (eBPF) via Cilium and Hubble, deriving the Communication Affinity Index (CAI): a novel composite heuristic that scores coupling through Transaction Amplification Rate, Network Overhead Ratio, and Cumulative Latency Penalty; components directly computed from the eBPF-collected Prometheus time-series. CAI scores drive three orchestration-layer remediations: pod colocation via GitOps-managed affinity rules, container merging, and a custom Kubernetes scheduler plugin (CAIScore). Evaluated on a standard microservice benchmark, CAIScore reduced median response times by 46.7%; container merging cut transaction latency by 73.2%; and GitOps-driven colocation lowered tail P_{95} by 55.6%, all without modifying application source code or disrupting developer workflows.

Keywords: *Microservices · Kubernetes · eBPF · Container Scheduling · Network Observability · Communication Affinity*

D-Track: Eliminating Redundant Fake Dirty Pages in Live Virtual Machine Migration

Samindu R. Cooray[†] , T. Kasun O. Mendis[†], Dinuni K. Fernando[*]

University of Colombo School of Computing, Sri Lanka

Live virtual machine (VM) migration facilitates the decoupling of the VM from its host, which is essential for fault tolerance, load balancing, and server maintenance. However, migration performance can degrade under memory-intensive workloads due to rapid memory page modifications and redundant page transfers. One key source of redundancy is fake dirty pages, where previously transferred pages are incorrectly marked as modified because of limitations in dirty page tracking. This paper proposes D-Track, an efficient dirty page tracking mechanism designed to mitigate the impact of fake dirty pages during pre-copy live migration. By reducing redundant page transfers, D-Track enhances migration performance by minimizing total migration time and downtime while preserving application performance. Our comprehensive evaluation on the KVM/QEMU platform demonstrates that D-Track, when combined with compression, reduces total migration time by 20% to 40%, with minimal overhead for applications running within the VM. When considering the downtime, network bandwidth degradation, and application overhead of the VM, the evaluation results depict that the D-Track performs in line with vanilla pre-copy without any noticeable degradation.

Keywords: *Fake dirty pages . Virtual Machines . Pre-copy migration.*

TECHNICAL SESSION 4

LLMs, Agentic AI & Trustworthy AI

Beyond Correctness: Output and Prompt Bloat in LLM-Generated Code

Pasandi Ranganath[1*], Chamath Keppitiyagama[1],Tharindu Wijethilake[1]

[1]University of Colombo School of Computing,Sri Lanka

While the evaluation of Large Language Model (LLM) code generation traditionally focuses on functional correctness, code that passes its tests can still be structurally inefficient. This output bloat is often compounded by prompt bloat, where the natural language prompt requires more tokens than the minimal reference code it generates. Does the verbal effort of specifying a minor task to an LLM then exceed the cost of writing the code directly? We study input and output bloat in Python code produced by GPT-5.5 Instant and Claude Sonnet 4.6 from natural prompts that eight individuals wrote for four domain-specific scenarios, accepting the first response, which gave sixty-three (63) analysable samples. Each was compared with a minimal reference solution using radon and vulture for static structure, tracemalloc and timeit for runtime cost, and functional correctness tests, and was manually annotated for six semantic bloat categories. The median generated code was 3.5 times the reference length, reaching 4.5 times in one scenario, while cyclomatic complexity grew far less (median ratio 1.0), because docstrings, example invocations and comments add lines without adding branches. Complexity-based metrics therefore miss much of the line-level bloat. Every prompt was longer than the reference code it described, and 54 of 61 generated functions were correct, with most failures traced to errors in the prompt.

Keywords*: Large Language Models, Code Generation, Code Bloat, Software Quality, Empirical Software Engineering.*

Reasoning Path Hijacking in Small Language Models: Vulnerability Assessment for Agentic AI Systems

Sandares Fernando[1*] , Randil Pushpananda[1], Ruvan Weerasinghe[2]

[1]University of Colombo School of Computing, Sri Lanka, [2]Informatics Institute of Technology, Sri Lanka

Agentic AI systems built on Small Language Models (SLMs) face Reasoning Path Hijacking (RPH), which is adversarial content that covertly redirects an agent's chain of thought toward attacker-controlled goals without exploiting any software vulnerability. To our knowledge, no benchmark combining these attack vectors, together with a detector comparison, exists for SLMs of at most 9B parameters. In this paper, we make three contributions. First, a 300-sample multi-vector benchmark spanning prompt injection, memory poisoning, and multi-turn stealth attacks. Second, we evaluate four SLMs (Qwen3.5-9B, Gemma-3-4B, LLaMA-3- 8B, Mistral-3B) under a unified ReAct framework, finding attack success rates of 40.9 – 64.4% (mean 52.1%). Third, we compare four detectors: a prompted Claude Sonnet 4 judge, a LoRA fine-tuned Qwen3.5-9B judge, Logistic Regression, and Random Forest. The fine-tuned judge achieves F1 = 0.872 at only +25.7% latency overhead but classical classifiers approach its precision at negligible latency with full interpretability. Our benchmark, detection pipeline, and practical detector selection guidance can be directly applicable to practitioners deploying agentic AI in security-sensitive environments.

Keywords: *agentic AI security · prompt injection · memory poisoning · reasoning path hijacking · small language models · LLM-as-a-Judge · ReAct.*

Knowledge Graph-Based Fact Inclusion Explanation for Abstractive Summarization

Nirasha Nelki*, Dilrukshi Gamage

University of Colombo School of Computing, Sri Lanka

Abstractive summarization models can generate factually unsupported content, while existing factual consistency methods mainly provide binary verification without showing how facts in the summaries relate to source evidence. We present KGFactExplainer, a post-hoc framework that generates source-grounded justifications for summary facts using constrained knowledge graph traversal. The framework decomposes source documents and summaries into subject-predicate-object triplets and constructs a knowledge graph from the source document's triplet list. A set of triplets is returned by the framework as the evidence paths to justify each fact triplet in the summary, along with metadata. A Qwen3-8B model is fine-tuned for evidence path justification retrieval. Experiments on a summarization model, PEGASUS-XSum outputs over a subset of the XSum dataset achieve evidence path agreement F1=0.719 against human-annotated evidence paths, improving 50.7% over the base model. Human evaluation shows that users perceive the generated evidence paths as useful and trustworthy.

Keywords: *Abstractive Summarization · Factual Consistency · Knowledge Graph · Post-hoc Explainability · Large Language Models.*

POSTER SESSION 1

Computer Vision & Graphics

Text2NeRF: Multi-Aspect Knowledge Distillation in Diffusion-Guided NeRFs for Text-Driven 3D Synthesis

Thisarani Wickramasinghe[1], RGN Meegama[2]

[1]NSBM Green University, Sri Lanka, [2]University of Sri Jayewardenepura, Sri Lanka.

Recent text-to-3D works unveil the effectiveness of diffusion-guided Neural Radiance Fields (NeRFs) in synthesizing three-dimensional content from natural language descriptions. However, semantic misalignments, view-inconsistent geometry, limited material realism, and restricted control of 3D representations have become critical challenges for existing methods. This work proposes Text2NeRF, a diffusion-guided NeRF framework for text-driven 3D synthesis, which is built upon an improved S3D-NeRF with a Semantic–Structural–Style Distillation, where CLIP-based semantic conditioning is coupled with Score Distillation Sampling, multi-view structural consistency constraints, and PBR-aware texture synthesis. In addition, Text2NeRF introduces attribute-level prompt control for fine-grained generation, multi-aspect 3D editing for selective text-driven modification, and automatic 3D quality refinement for iterative improvement of representations. Structural supervision is enforced with RGB, depth, and normal consistency, while PBR-aware style distillation decouples material refinement from geometry optimization. Experiments on ShapeNet and Objaverse evaluate the framework using geometric, semantic, perceptual, and structural metrics. The results unveil the effectiveness of jointly optimizing semantic alignment, structural coherence, material realism, controllability, and iterative refinement
for text-driven 3D synthesis.

Keywords: *Diffusion Models, Neural Radiance Fields, Text-to-3D Generation, Score Distillation Sampling, Knowledge Distillation, Controllable 3D Generation, 3D Editing, Physically Based Rendering*

Deployment-Aware Optimization of Real-Time Neural Rendering Under Fixed Runtime Constraints

B.H.S. Tharuka Fernando[1*], P.P.G. Dinesh Asanka[1]

Department of Industrial Management, University of Kelaniya, Sri Lanka

Real-time neural rendering performance depends on model accuracy, exportability, backend numerical stability, and runtime-path overhead. This paper studies those factors jointly in a geometry-guided neural rendering pipeline. The workflow combines candidate screening, backend-validated mixed-precision deployment, and profiling across S1 (FastAPI HTTP), S2 (native ONNX Runtime), and S3 (TensorRT zero-copy). Experiments use 7,000 aligned triplets at 704 × 704 with batch-1 inference on an RTX 5070. Candidate screening selected DepthNet-EE for geometry support and HybridStereoNet-FP16 for rendering. Backend validation showed that export success alone was insufficient: PyTorch-to-ONNX drift stayed at 0.2% mean absolute error, while the first ONNX-to-TensorRT build reached 4.9% before correction. The reported S1 stage summary was 142 ms; the S3 Unity overlay displayed 21.69 ms for a processed frame and a rolling average of 47.1 FPS. The final precision-constrained TensorRT engine reached 37.09 dB PSNR and 0.9655 SSIM.

Keywords: *real-time rendering, deployment optimization, mixed precision, ONNX Runtime, TensorRT, zero-copy execution*

CloneVerify: Quality-Aware Conditional Enhancement for Fine-Grained Sri Lankan Tea Clone Identification

Sadesh Weerasuriya[1]*, Imesh Pathirana[1]
[1]*Department of Computing, Informatics Institute of Technology,Sri Lanka*

Sri Lankan tea clone identification depends on scarce specialist expertise, and automated recognition is difficult because cultivars are morphologically similar and field images are uncontrolled. This paper presents CloneVerify, a quality-aware pipeline for fine-grained identification of six Tea Research Institute (TRI) clones from field-captured RGB images. It combines a multi-task Image Quality Gate that predicts PASS/ENHANCE routing together with blur, noise, detail-loss, and exposure scores; LeafSRNet ×4, a tea-leaf super-resolution network with a Leaf Structure Attention Module and a two-stage upsampler; and a Hybrid SimCLR classifier that combines contrastive pretraining with supervised fine-tuning. A TRI-verified dataset of 1,620 images across six clones was divided using a stratified 70/15/15 protocol. The classifier achieved 83.26% test accuracy (95% CI 77.7–87.7%) and 0.8396 macro F1 on 215 unseen field images. LeafSRNet improved over bicubic ×4 by 1.46 dB PSNR and 0.0581 SSIM. The quality gate routed every degraded sample to enhancement at Tpass = 0.80; its ROC-AUC of 0.5645 and high enhancement rate on clean images indicate a conservative operating point calibrated to synthetic degradations. These results support the feasibility of conditional enhancement for field-oriented clone recognition and identify multi-site data, real degradations, and an external test set as the next steps.

Keywords: *Tea clone identification, quality-aware routing, super resolution, SimCLR, fine-grained classification, agricultural computer vision*

Small Object Detection Enhancement in Real-Time Detection Transformer (RT-DETR) for UAV Imagery

J.M. Sasanka Madhura Bandara [*] , Pumudu A. Fernando

School of Computing, Informatics Institute of Technology, Sri Lanka

Small object detection (SOD) in UAV aerial imagery is one of the most persistent unsolved challenges in computer vision, characterised by severe scale variation, dense object clustering, frequent occlusion, limited pixel representation (objects <32×32 px), and complex backgrounds that collectively degrade detector performance. The Real-Time Detection Transformer (RT-DETR) has established a new paradigm by surpassing YOLO models in both accuracy and real-time speed on standard benchmarks, achieving 53.1 mAP on COCO at 108 FPS with a ResNet-50 backbone. However, it struggles to reliably localise small or tightly grouped targets in challenging datasets such as VisDrone2019. This paper proposes targeted model-level enhancements to the RT-DETR architecture: integration of a P2 feature pyramid level, Space-to-Depth (SPDConv) convolution for information-preserving downsampling, and a novel Coordinate-Aware Gated Convolution (CoordGnConv) module for spatially sensitive feature fusion. Evaluated on VisDrone2019, the proposed model achieves a mAP@50 of 44.1% (+1.8%), mAPsmall of 18.1% (+1.9%), and mAP@50–95 of 26.3% (+1.4%) over the baseline, while maintaining real-time inference at 86 FPS with 19.2M parameters. A comprehensive ablation study and cross-GPU benchmarking across five hardware platforms confirm the complementary nature of the three modules and the scalability of the proposed design.

Keywords: *RT-DETR, Small object detection, UAV imagery.*

POSTER SESSION 2

Intelligent Sensing, IoT & Adaptive Learning

Towards Non-Visual Olfactory Sensing: An Intelligent Electronic Nose for Detecting Sri Lankan Elephants (Elephas maximus maximus)

W.M.A.T Wanninayake[1] , P.W Paranawithanage[1] , T.S Dissanayake[1], W.M.H.K Weerakoon[1], R.G.B.C Rajagalgoda[1], B.M.G.S.T.S.K Wickramasinghe[1 *], N.A.A.M Prabhath[1], G.D Illeperuma[1], T.S.P Fernando[1]

[1] *The Open University of Sri Lanka, Nawala, Nugegoda*

The Asian Sri Lankan elephant (Elephas maximus maximus) is endangered in Sri Lanka and faces growing threats from habitat fragmentation, urbanization, and rising Human-Elephant Conflict (HEC). Traditional monitoring tools like camera traps are often expensive, infrastructure-dependent, and impractical in remote areas. Chemical ecology research shows elephants emit distinct Volatile Organic Compounds (VOCs) through their temporal glands, urine, skin, and breath, offering a non-visual, non-invasive detection alternative. This paper presents a novel, intelligent electronic nose designed to detect Sri Lankan elephants. The prototype features a multi-sensor array of semiconductor gas sensors, configured to detect and characterize variations in VOC-related gas signatures associated with elephant presence. The device utilizes a microcontroller that works as the primary processing unit for real-time processing. To address the complexities of field-based olfactory sensing, the e-nose runs on-device machine learning for autonomous classification of scent profiles, and its architecture incorporates advanced sensor drift compensation and environmental noise filtering to ensure high sensitivity. Technical evaluations of the prototype focus on the classification accuracy using a novel field dataset, hardware practicability, and the limit of detection. The embedded Random Forest algorithm achieved 95.8% classification accuracy in independent testing, and field validation confirmed a reliable detection range of up to 10 meters. This e-nose offers a non-invasive alternative to traditional wildlife monitoring technologies, providing a scalable tool for scent detection. The research concludes by highlighting current results and limitations in VOC data availability, sensor reliability, and field validation, and proposes future research directions for developing robust odor-based wildlife monitoring platforms.

Keywords: *VOC detection, Olfactory sensing, Odor signature, Elephant conservation.*

Quantum-Aware Adversarial Defense for IoT Intrusion Detection Systems Using Explainable AI

Supun P. Shaminda[1,*], Charaka Geethal Kapugama[1]

[1] *Department of Computer Science, Faculty of Science, University of Ruhuna, Sri Lanka*

The Internet of Things (IoT) has placed billions of constrained devices on public networks, and the machine-learning intrusion detection systems (IDS) that protect them can be fooled by small adversarial perturbations. Quantum search is often cited as a future amplifier of this threat, yet the real cost of such an attack is rarely examined. This paper presents a quantum-aware and explainable defence framework for IoT IDS. A lightweight one-dimensional convolutional neural network with 1,106 parameters is trained on ToN_IoT with on-the-fly Fast Gradient Sign Method (FGSM) samples and evaluated against FGSM at three perturbation budgets. Vulnerability density is measured with the Grover Vulnerability Rate (GVR) on real attack traffic and on a stress set produced by a Quantum Circuit Born Machine (QCBM). The Grover analysis is then made concrete: we define the search space and a reversible oracle, give the amplitude-amplification circuit and estimate its fault-tolerant cost. SHAP explanations make every alert readable for an analyst. The detector reaches 92.5% clean accuracy, keeps 91.1% under single-step FGSM at $\varepsilon = 0.10$, and shows a real-traffic GVR of 0.85% (95% CI 0.79–0.92%). One coherent oracle call needs about 7.7×105 Toffoli gates and 6×103 logical qubits, so Grover-style probing beats classical random probing only when fewer than about one input in 109 is misclassified. GVR is therefore a quantum-aware robustness indicator, not evidence of resistance to a real quantum attacker.

Keywords: *IoT security, Intrusion detection, Adversarial robustness, Explainable AI, Grover search, Quantum resource estimation, Quantum circuit Born machine.*

SmartFedProx: Personal Finance Modelling Using Federated Learning to Address Data Heterogeneity

V A D Jeremy Perera[1 *], Torin Wirasingha[2]

[1] *Department of Computing, Informatics Institute of Technology,Sri Lanka,*

[2] *Department of Computing, Informatics Institute of Technology, Sri Lanka*

Data privacy in machine learning, handling non-independent and identically distributed (non-IID) data across financial institutions and preventing model divergence in decentralised training are three main challenges identified that traditional personal finance modelling faces. This research proposes SmartFedProx, an enhanced federated learning framework extending traditional centralised machine learning solutions for personal finance regression. The framework addresses data heterogeneity through an approach combining, (1) Two level adaptive proximal regularisation that dynamically updates the coefficient μ at the client level based on client divergence from the global model and on the server level based on the global loss trends and (2) Hybrid Client Selection Strategy that balances a selection of High divergence clients, stable middle divergence clients and low divergence clients. Across all simulation configurations, SmartFedProx achieves higher R^2, lower MSE, and faster convergence than both baselines. The R^2 reaches 0.7571 within 25 rounds, compared to 0.6038 for FedProx and 0.4701 for FedAvg under identical conditions.

Keywords: *Federated Learning, Personal Finance Modelling, Data Heterogeneity, Adaptive Proximal Regularisation, Client Selection*

A Hybrid Concept Drift-Aware Streaming Learning for Real-Time IT Project Risk Prediction

W. K. C. P. Bandara*, A. M. C. Malkanthi

Department of Physical Sciences and Technology, Faculty of Applied Sciences, Sabaragamuwa University of Sri Lanka.

IT project data streams are highly dynamic and subject to concept drift. The statistical relationships between issue attributes and risk labels continuously evolve due to changes in team structures, development methodologies, and project environments over time. Static batch models trained on historical snapshots may degrade between retraining cycles and are therefore ill-suited for real-time deployment. This paper presents a concept drift-aware streaming learning framework for real-time IT project risk prediction using Apache JIRA issues from Apache Software Foundation open-source projects. A hybrid batch-streaming architecture is introduced by incorporating the posterior probability generated by a pretrained LightGBM model as an additional feature for streaming learners. Seven streaming classifiers are evaluated against a LightGBM batch baseline under the prequential evaluation protocol. The results indicate that concept drift is a recurring characteristic of the data stream, with ADWIN and DDM detecting 54 and 56 drift events, respectively, across the 189,282-step test stream. Despite operating in a single pass without full dataset storage or periodic retraining, Hoeffding Tree (F1 = 0.9416), Incremental MLP (F1 = 0.9410), and ADWIN Bagging (F1 = 0.9398) achieved predictive performance comparable to the LightGBM batch baseline (F1 = 0.9421). Further analysis confirms that the observed performance gains arise from the complementary strengths of batch-derived prior knowledge and adaptive streaming learning. These findings demonstrate that both batch-derived prior knowledge and adaptive streaming learning contribute to predictive performance, supporting the effectiveness of the proposed hybrid framework for real-time IT project risk prediction under concept drift conditions.

Keywords: *Concept drift, Streaming machine learning, IT project risk prediction, Hybrid learning, Apache JIRA.*

POSTER SESSION 3

Applied Data Analytics and Machine Learning

SinFair: Mitigating Demographic Entity Erasure in Sinhala Abstractive Summarization via Five-Axis Counterfactual Augmentation

Anupa Wanigasinghe[1*], Salitha Perera[1]
[1] *School of Computing, Informatics Institute of Technology, Sri Lanka*

Sinhala is a low-resource Indo-Aryan language spoken by approximately 16 million people, yet existing Sinhala abstractive summarization systems do not address demographic bias. This paper proposes SinFair, targeting a previously unmeasured failure mode termed Entity Erasure, where abstractive models achieve apparently low semantic disparity by deleting demographic subjects rather than retaining them. To detect this, ERR (Entity Retention Rate) is introduced, a metric applying a grammatical stop-root filter to isolate demographic noun retention. A five-axis counterfactual dataset of 63,210 samples is constructed from 9,030 NSina news articles, covering gender, honorifics, age, social status, and regional affiliation. SinFair is trained using a two-stage curriculum: Stage 1 establishes Sinhala summarization capability, while Stage 2 fine-tunes on the counterfactual dataset under a constrained learning rate prioritizing ROUGE-L preservation. The fine-tuned mT5-base improved overall ERR from 45.97% to 48.25%, and mBART-large-50 from 46.10% to 50.16%, with honorifics showing the largest gains in both models (+6.60pp mT5-base, +10.48pp mBART), though only the mBART improvement reached statistical significance. Dimensional analysis reveals a retention hierarchy: gender markers show high retention, regional names reach a factual ceiling, honorifics yield the largest learning signal, and age markers are actively resisted by the compression objective. Zero-shot evaluation of XL-Sum mT5 and Aya-101, both achieving ERR below 20%, is consistent with Entity Erasure extending to large-scale multilingual systems, suggesting scale alone is insufficient to resolve demographic entity retention in native Sinhala summarization.

Keywords: *Sinhala NLP, abstractive summarization, demographic Entity Erasure, counterfactual augmentation, Entity Retention Rate, honorifics, low-resource language.*

Federated Reasoning: A Privacy-Preserving Medical Question Answering System with Chain-of-Thought Reasoning and Retrieval-Augmented Generation

J. G. Sevindu Imsara[1,2], Janith Prabhanuka[1*]

[1] *School of Computing, Informatics Institute of Technology, Colombo, Sri Lanka* [2] *Robert Gordon University, Aberdeen, Scotland*

Deploying large language models in clinical environments faces two persistent barriers: regulatory constraints that prohibit centralised data aggregation and the lack of transparent reasoning required for clinical trust. Building on recent work on federated chain-of-thought and federated LoRA, this paper presents a controlled empirical comparison of centralised and federated chain-of-thought medical question answering using parameter-efficient fine-tuning and retrieval augmentation. An 8-billion-parameter Meditron model is fine-tuned via LoRA adapters within a three-client Flower simulation under non-IID Dirichlet partitioning. A hybrid dense-BM25 retrieval module over a 7838-passage biomedical corpus is integrated at inference. Across open-ended clinical QA, the ten-round federated model matches and marginally outperforms a centralised baseline on BLEU, ROUGE, and BERTScore metrics, demonstrating that privacy-preserving training does not inherently degrade generative reasoning quality in this setting. A task-specific trade-off is observed on MedQA-USMLE, where chain-of-thought optimization reduces structured factual recall. These results suggest that the privacy-utility trade-off in federated medical NLP is governed more by training depth than architectural limitation.

Keywords: *Federated Learning, Medical Question Answering, Chain-of-Thought Reasoning, Retrieval-Augmented Generation, Privacy-Preserving AI, Large Language Models*

Quantifying Metformin Confounding in Machine Learning-Based Metagenomic Glycemic Classification

Ruchila Wijerathne[1], Rupika Wijesinghe[1*]

[1] *University of Colombo School of Computing, Sri Lanka*

While published Type 2 Diabetes (T2D) microbiome classifiers report high accuracy (AUC 0.71–0.85), the contribution of metformin confounding remains unquantified at the pipeline level. The objectives of this study are to systematically quantify the confounding effect of metformin across multiple machine learning classifiers, isolate robust, medication-independent microbial biomarkers using SHAP-based Explainable AI, and evaluate the classification boundaries characterizing early-stage pre-diabetes (Control vs. IGT) and subsequent disease progression (IGT vs. T2D). We present a systematic evaluation using the European MetaCardis cohort (n=1,467), spanning six disease-stage comparisons with and without metformin exclusion. Using seven tuned machine learning architectures, we demonstrate that diagnostic performance decreases substantially upon metformin exclusion: Control vs. T2D AUC drops from 0.844 to 0.672, while IGT vs. T2D progression classification falls from 0.807 to 0.604, where wide confidence intervals reflect limited statistical power. Control vs. IGT detection remains detectable but weak (AUC 0.681). SHAP analysis isolated robust, medication-independent biomarkers, including Clostridium sp. CAG:167, Agathobaculum butyriciproducens, and Parasutterella excrementihominis, that persist as discriminative features in drug-naive models. These results provide the first pipeline-level quantification of metformin confounding, proving that while published performance is largely medication-driven, a genuine but attenuated biological signal characterizes drug-naive T2D.

Keywords: *gut microbiome, Type 2 Diabetes, Metformin confounding, machine learning, SHAP, curatedMetagenomicData*

Attitudinal Predictors of Music Genre Exploration and Abandonment Among Streaming Users: A Survey-Based Multi-method Analysis

Weerasuriya Arachchige Malsha Prabuddhi[1*], Saminda Chandika Premaratne[2], Podduwa Kankanamage Subash Chaminda Jayasinghe[1]

[1] Department of Information and Communication Technology, Faculty of Technology, University of Ruhuna, Sri Lanka. [2] Department of Information Technology, University of Moratuwa, Sri Lanka

This paper reports a cross-sectional self-report survey exploring genre-oriented music listening behavior among music streaming listeners, considering psychological and attitudinal drivers of users' genre preferences, exploration, and abandonment. Using a sample of n = 156, the study examines listening behavior with exploratory factor analysis (EFA), K-means clustering, ordinary least-squares (OLS), logistic and proportional-odds regression, chi-square tests, and Spearman correlations, with bootstrap 95% confidence intervals (B = 1,000 unless stated otherwise). First, nostalgia orientation was the strongest positive correlate of self-reported genre exploration orientation (B = 0.291, $p < .001$), and the association remained after adjusting for response style. Second, listeners were partitioned into two engagement segments, High-Engagement (n = 95, 60.9%) and Passive (n = 61, 39.1%), which were stable but weakly separated (silhouette = 0.232). Third, genre diversity and artist loyalty were additional positive predictors of exploration orientation (B = 0.264, p = .023, and B = 0.189, p = .019), although the loyalty effect was not robust to response-style adjustment. Finally, poor audio quality was the most frequently cited potential reason for abandoning a genre (65.4%). EFA identified three dimensions, Intrinsic Song Qualities ($\alpha = 0.811$), Curation & Discovery ($\alpha = 0.809$) and Listener Attitude ($\alpha = 0.647$), with good sampling adequacy (KMO = 0.838). Limitations include the predominantly Sri Lankan convenience sample and the dominance of classical music listeners. Consequently, this paper offers candidate attitudinal variables for future testing against track-level streaming data.

Keywords: *Music Genre Preference, Streaming Behavior, Exploratory Factor Analysis, Genre Exploration, Genre Abandonment, Bootstrap Confidence Intervals, Listener Segmentation.*

POSTER SESSION 4

Software and Systems

Enhancing Live VM Migration for Latency-Sensitive Workloads

Ayash Maddumage[1] , Nadeesha Epa[1] , Dinuni K. Fernando[1*]

University of Colombo School of Computing, Colombo, Sri Lanka

Live virtual machine (VM) migration is widely used in cloud environments for maintenance, load balancing, fault recovery, and resource consolidation. However, latency-sensitive workloads - such as video streaming, online gaming, and interactive database services - can experience severe disruption during migration even when traditional hypervisor-level metrics appear acceptable. This paper investigates live VM migration from both infrastructure and application perspectives in a QEMU/KVM testbed across five representative workloads, three migration families (Precopy, Postcopy, and Hybrid), and five optimization settings. Our empirical results verify a pronounced divergence between hypervisor-level metrics and application-level behavior: configurations minimizing Total Migration Time (e.g., Precopy with Multi-FD) or hypervisor-reported downtime (e.g., Postcopy and Hybrid achieving sub-15 ms downtime) frequently induce multi-second application stalls. Across all workloads, Precopy-based configurations delivered the most stable application responsiveness, with XBZRLE and Compression emerging as the optimal workload-dependent optimizations. Selecting the empirically optimal workload-specific Precopy configuration reduced mean peak application disruption by 43.3% on average relative to vanilla Precopy, outperforming static policies of fixed Compression (34.5%) and fixed XBZRLE (20.2%). Motivated by these findings, we present a lightweight QMP-based adaptive migration framework that performs short pre-migration dirty-memory profiling to dynamically select between XBZRLE and Compression, trading a bounded, distributed pre-migration profiling disturbance for a substantial reduction in main-migration service interruption.

Keywords: *Live VM migration · Latency-sensitive workloads· Application-level disruption · Precopy migration · QEMU/KVM · Adaptive migration framework ·XBZRLE · Compression*

When the Next Step Is Not One Step: Distribution-Aware Execution Modeling for Concurrent Go Programs

Kaviru Hapuarachchi[1]

[1]*University of Colombo School of Computing, Colombo, Sri Lanka*

Predicting the next scheduler event of a concurrent program is ill-posed as single-label classification: repeated runs from the same trace prefix produce different, equally valid next events. We formulate next-event prediction for concurrent Go programs as estimating an empirical next-event distribution obtained from repeated executions, and release a benchmark with 945 training examples from hand-written and generated programs and 798 held-out predictions on 54 real-world GoKer bug kernels from systems such as CockroachDB, Kubernetes, gRPC and etcd. We fine-tune Qwen2.5-Coder-7B with low-rank adapters under a cross-entropy objective and a KL-to-empirical-distribution objective and compare them with zero-shot models and trivial baselines. The outcome is a neutral-to-negative result. The fine-tuned adapters (re-evaluated: 36.3% cross-entropy, 33.6% KL; originally reported 36.2% and 35.8%) are statistically indistinguishable from always predicting the majority event (35.5%), and an event n-gram model reaches 33–34%. Because each heldout prompt carries labels from about five runs, no deterministic predictor can exceed 58.9% top-1 accuracy, so we report macro-F1, balanced accuracy, per-class recall, program-clustered confidence intervals and binned calibration measured directly on both adapters. A sensitivity study with 100 runs per program shows that five-run targets are noisy (mean total-variation distance 0.24 from a 100-run reference), which plausibly limits what the KL objective can learn. Rare event types are essentially never predicted. A select-block goroutine-leak signature follows from scheduler semantics rather than learning. We release the data, adapters and tooling.

Keywords: *Concurrent execution modeling · Nondeterminism · Distribution estimation · Calibration · Go · Concurrency bugs.*

INDUSTRY TRACK

Development of an Intelligent and Secure Smart Home Automation Using Home Assistant and IoT Devices

[1]Rumesh Rajasegar, [2]Thayaparan Subramaniam, [2]Mukunthan Tharmakulasingam,

Sri Lanka Telecom PLC, Sri Lanka, irumesh.work@gmail.com, University of Moratuwa, Sri Lanka, thayaparan@uom.lk, University of Jaffna, Sri Lanka, mukunthan@eng.jfn.ac.lk

Keywords: IoT, Home Automation, Home Assistant, Voice Assistant, STT, TTS, LLM, Smart Home Security

Abstract: The growing integration of Internet of Things (IoT) devices into daily life has transformed modern homes into connected, automated environments. However, fragmented multi-vendor ecosystems, inadequate security, and limited device interoperability remain key barriers to widespread smart home adoption. This paper presents the development of an intelligent and secure smart home automation system built on Home Assistant (HASS), an open-source platform ranked first among 20 comparable systems. The proposed system integrates diverse sensors, actuators, IP cameras, and custom ESP32 microcontrollers, complemented by voice assistant pipelines using local speech-to-text and text-to-speech models, and large language model (LLM) integration for natural-language device control. The system demonstrates unified control across heterogeneous devices from multiple vendors, remote WAN access without cloud dependency, presence-based and scheduled automation, and privacy-preserving local data processing. Voice command latency under 2 seconds was achieved using local STT/TTS pipelines. Results confirm the viability of open-source, locally operated smart home systems as a secure, flexible, and cost-effective alternative to proprietary solutions. Future work will address redundancy mechanisms, adaptive AI-driven automation, and simplified onboarding for non-technical users.

1 Introduction

The rapid proliferation of Internet of Things (IoT) devices has fundamentally transformed modern residential environments into connected, data-driven ecosystems. Taiwo et al. defined a taxonomy of home automation systems across five components: application area, automation layers, protocols, platforms, and sensors [4]. Dong et al. traced the evolution from self-powered sensors toward AIoT-enabled smart homes, highlighting the convergence of AI and IoT as the next frontier [8]. Despite rapid growth, widespread

adoption is hindered by fragmented multi-vendor ecosystems, inadequate security, and poor device interoperability. Yar et al. demonstrated that edge-computing paradigms significantly improve IoT smart home responsiveness and privacy [5]. Open-source platforms offer cost-effective solutions: Home Assistant (HASS), ranked first across 20 open-source home automation systems by Setz et al. [1], now supports different vendor devices as bundled as integrations in Home Assistant [2]. This paper presents a HASS-based system which integrates diverse sensors, actuators, cameras, voice assistants, and large language models (LLMs) to deliver unified, privacy-preserving smart home automation.

2 Methodology and System Design

The system was deployed on a Raspberry Pi 4B (4 GB RAM) running Home Assistant OS, chosen for local-first processing and full data sovereignty at low cost ($55). Security was a primary design concern; Allifah and Zualkernan ranked IoT smart home device security and identified local-processing architectures as preferable to cloud-dependent ones [6]. Hardware integration spans three categories:

(i) **Sensors** – PIR, temperature, humidity, CO_2, air quality, gas leak, contact, voltage, and water leak sensors;

(ii) **Actuators** – smart switches, bulbs, relay controllers (EMR/SSR), curtain sliders, and door locks; and

(iii) **Cameras** – integrated via ONVIF/RTSP protocols.

Inspired by Stolojescu-Crisan et al., who built a smart home system using ESP microcontrollers [7], a custom ESP32 module with a PIR sensor and relay was developed via ESPHome to demonstrate motion-triggered lighting automation.

Remote WAN access was configured through ISP port-forwarding without cloud subscriptions. Automations follow a Trigger–Condition–Action model covering presence-based lighting, scheduled climate control, and security arming on departure. Cross-vendor automation was demonstrated by triggering Tuya switches, with action from WolfGuard Alarm system via HASS.

Voice interaction used an M5 Atom Echo satellite with a custom wake word trained via Google Colab. Faster-Whisper [3] provided local speech-to-text (STT) and Piper handled local text-to-speech (TTS), with Microsoft Azure cloud STT/TTS as a higher-accuracy

fallback. LLM integration via the OpenAI Conversation integration (HASS 2024.10) enables natural-language device control, with carefully scoped entity exposure to mitigate hallucination risks. Fig 1 illustrates the full System Architecture.

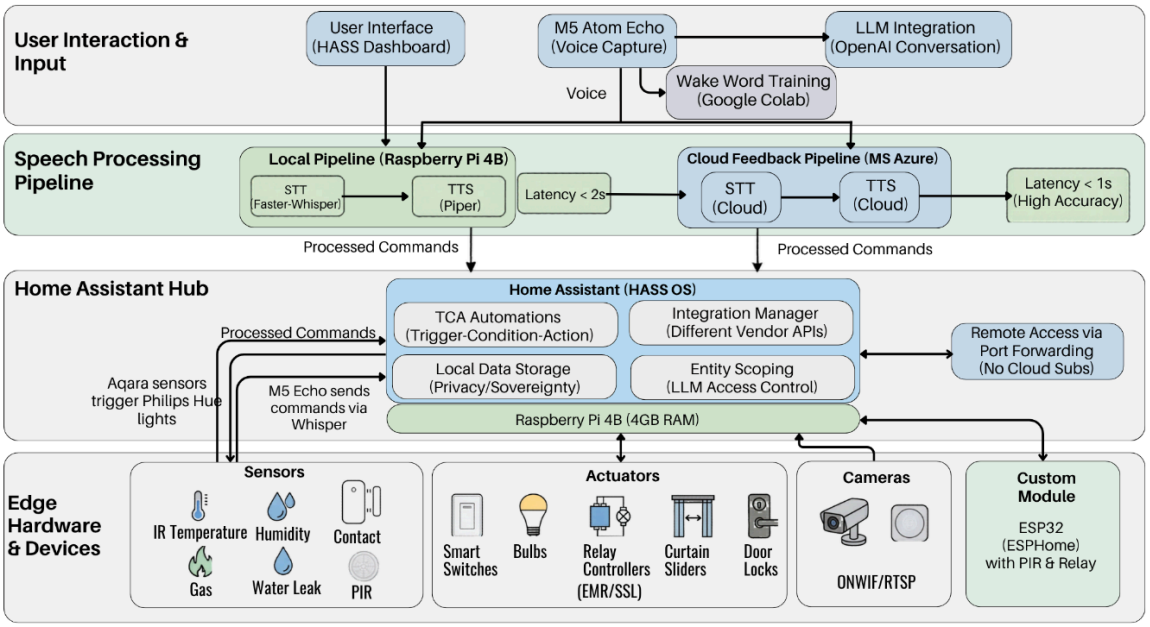


Fig. 1. System Architecture

3 Results and Discussion

The system achieved unified, reliable control across heterogeneous devices from vendors including Tuya, WolfGuard, Philips Hue, Aqara, Hikvision, and Shelly through a single dashboard, eliminating the need for separate vendor applications. Presence-based automation reduced unnecessary lighting and HVAC operation. Consistent with the edge-computing findings of Yar et al. [5], local STT/TTS processing yielded voice command latency under 2 seconds, while cloud-based processing delivered sub-second response with improved accuracy. LLM-assisted queries enabled non-technical users to interrogate device states conversationally. Aligned with the security framework of Allifah and Zualkernan [6], all sensor data was processed and stored locally, ensuring full privacy without third-party cloud dependency.

Identified limitations include higher electricity consumption at scale, vulnerability during power or network outages due to absent redundancy, and a configuration learning curve for non-technical users. LLM hallucination risk requires deliberate scoping of which entities are exposed to AI agents.

4 Conclusion

This research demonstrates that Home Assistant, as established by Setz et al. [1] to be the leading open-source home automation platform, augmented with DIY ESP32 microcontrollers, voice AI pipelines, and LLMs, provides a secure, flexible, and cost-effective foundation for intelligent smart home automation. Addressing the taxonomy defined by Taiwo et al. [4] across automation layers, protocols, and sensors, the system unified heterogeneous devices, enabled privacy-preserving voice interaction, and demonstrated LLM-based accessibility. Extending the AIoT vision of Dong et al. [8], future work will target adaptive AI-driven automation learning from user behaviour, battery-backup redundancy, energy-efficient edge computing, and simplified onboarding interfaces. This work contributes practical insights toward scalable, privacy-preserving smart home systems aligned with ICTer's focus on impactful ICT applications for societal benefit.

References

1. Setz, B., Graef, S., Ivanova, D., Tiessen, A., Aiello, M.: A Comparison of Open-Source Home Automation Systems. IEEE Access 9, 167332–167352 (2021). doi: 10.1109/ACCESS.2021.3136025
2. Home Assistant: Integrations. https://www.home-assistant.io/integrations/ (Accessed Jun 2026)
3. SYSTRAN: faster-whisper. https://github.com/SYSTRAN/faster-whisper (2024)
4. Taiwo, O., Gabralla, L.A., Ezugwu, A.E.: Smart Home Automation: Taxonomy, Composition, Challenges and Future Direction. In: ICCSA 2020, LNCS, vol. 12254, pp. 878–894. Springer (2020)
5. Yar, H., Imran, A.S., Khan, Z.A., Sajjad, M., Kastrati, Z.: Towards Smart Home Automation Using IoT-Enabled Edge-Computing Paradigm. Sensors 21(14), 4932 (2021)
6. Allifah, N.M., Zualkernan, I.A.: Ranking Security of IoT-Based Smart Home Consumer Devices. IEEE Access 10, 18352–18369 (2022)
7. Stolojescu-Crisan, C., Crisan, C., Butunoi, B.-P.: An IoT-Based Smart Home Automation System. Sensors 21(11), 3784 (2021)
8. Dong, B., Shi, Q., Yang, Y., Wen, F., Zhang, Z., Lee, C.: Technology evolution from self-powered sensors to AIoT enabled smart homes. Nano Energy 79, 105414 (2021)

Quantum Risk Assessment Framework for Harvest Now, Decrypt Later (HNDL) Attacks in Banking Sector

Vishvadini Kurukulasooriya*[1], S. Kasun Chathuranga[1], Udara Wijesinghe[2]

[1]*TechCERT, Colombo, Sri Lanka, Department of Computer Science and Engineering,*
[2]*University of Moratuwa, Sri Lanka*

The banking industry relies on conventional cryptographic mechanisms including asymmetric key cryptosystems such as Rivest-Shamir-Adleman (RSA), Elliptic Curve Cryptography (ECC) and Diffie-Hellman as well as symmetric key cryptosystems such as the Advanced Encryption Standard (AES) to secure bank transactions, interbank communications, payment systems and other digital banking functions. Recent advancements in quantum computing have demonstrated that Cryptographically Relevant Quantum Computer (CRQC) could undermine widely deployed public key cryptographic schemes and reduce the security margin of symmetric cryptography. In particular, Harvest Now Decrypt Later (HNDL) attacks allow adversaries to intercept and store encrypted information for potential decryption when sufficiently capable quantum computing become available. While existing PQC guidance provided by the NSA, NIST and the European Commission focuses primarily on long term migration strategies and targets, a methodology to assess and prioritize HNDL exposure in a banking sector environment remains insufficiently addressed. In this paper we propose a framework for quantum risk assessment against Harvest Now, Decrypt Later (HNDL) attacks in the banking sector environment and presents a scoring model to prioritize quantum risk of external communication channels based on data confidentiality, data longevity, business impact and likelihood of being targeted for interception. By providing a remediation roadmap to support post quantum migration strategies for identified external communication channels, we focus to reduce the current HNDL attack exposure of confidential banking data.

Keywords: *quantum-risk-assessment · harvest-now-decrypt-later-attacks · post-quantum-cryptography · banking-data · confidentiality*

An AI-Enabled Digital Learning Passport with Skill Gap Detection and Industry-Integrated Recruitment Interface: A Design Science Approach

Madusanka Premaratne Rathnayake Mudiyanselage[1*] , Hasanthi Lakmali Thellapura Arachchilage[1] , Dimuthma Umashani[1] ,Dillina Herath[2]

[1]Knivok Private Limited, Sri Lanka
**madusanka@knovik.com, hasanthi@knovik.com, dimuthma@knovik.com*
[2] ESOFT University, Sri Lanka
dillina@esu.lk

Learning management systems record detailed evidence about how students learn, but the record that follows graduates into employment is still a static transcript. Students have no control over that record and lose access to it on graduation; recruiters, in turn, cannot verify the competence claims that replace it. This paper applies Design Science Research to develop a Digital Learning Passport (DLP): a learner-controlled record that turns raw LMS signals into a portable skill record designed for offline verification by recruiters. We demonstrate the engine on the Open University Learning Analytics Dataset, constructing 29,228 complete registration-level profiles. Aggregate gap severity correlates with held-out outcomes (Pearson r = −0.767), and severity-ranked screening identifies at-risk outcomes with F1 = 0.881. A recruitment process model shows the largest effort reductions fall at CV screening and first-stage competency checks.

Keywords: *Digital Learning Passport · Learning analytics · Skill gap detection · OULAD · Skills-based hiring · Design science research.*

An Edge-Cloud Machine Condition Monitoring System for Industrial Rubber Processing

Rumesh Rajasegar
Research and Development Section, Sri Lanka Telecom PLC, Sri Lanka

Rubber compounding operations depend on the reliable performance of Internal Mixers and Two-Roll Mills, yet continuous machine monitoring remains impractical for most facilities due to the high cost, integration complexity, and limited connectivity of conventional PLC-based systems. This paper presents an edge-cloud machine condition monitoring system that addresses this gap through three contributions. First, a compact ESP32-based embedded gateway integrates two current transformers (600/5A), one infrared temperature sensor (0–200°C, 4–20mA), and one K-type thermocouple (0–400°C, 0–10V) within a single unit, eliminating the need for multiple disparate monitoring devices. Second, an edge processing pipeline performs signal conditioning such as ADC conversion, RMS computation, spike filtering, and sensor health detection, transmitting processed data at one-second intervals via MQTT over 4G LTE to a cloud dashboard, enabling root cause analysis of machine incidents and predictive maintenance planning. Third, a Plug and Play provisioning model eliminates device onboarding complexity, enabling rapid deployment across active production lines. Validated across four live production units on Internal Mixers and Two-Roll Mills, the deployed system demonstrates reliable detection of overload, overheat, no-load, and idle machine conditions, providing a practical and scalable path from reactive to predictive maintenance in rubber compounding operations.

Keywords: *Machine Condition Monitoring, IIoT, MQTT, FreeRTOS, ADC, Symmetric Multiprocessing, Rubber Compounding, Predictive Maintenance*

ICTer 2026

26th International Conference
on Advances in ICT for Emerging Regions

Organized by

UCSC University of Colombo
School of Computing

SPONSORS

University of Colombo School of Computing
35, Reid Avenue, Colombo 7
SRI LANKA

https://ucsc.cmb.ac.lk/