



IEEE
Industrial
Electronics
Society



IEEE INTERNATIONAL CONFERENCE ON RESPONSIBLE AI

La Trobe University
2 – 5 September 2026

PROGRAM

Wednesday 2 September 2026		
09:00-09:30	📍 Lunch/Coffee Area Reception Desk + Breakfast	Coffee
09:15-09:30	📍 Lunch/Coffee Area Welcome to IRAI	Plenary
09:30-12:00	📍 301 AI Skills: Prompting and Vibe Coding for Industry Practice AI Skills: Prompting and Vibe Coding for Industry Practice A practical session on turning workplace problems into useful AI workflows, from better prompting to rapid AI-assisted prototyping. Participants learn reusable prompt patterns, how to brief AI coding tools, and how to move from idea to a lightweight internal tool while managing reliability, privacy and review. Nishan Mills is a senior AI architect, data scientist, educator and entrepreneur with more than 15 years of experience addressing technology challenges across multinational, multisector digital environments including healthcare, emergency services, transportation and sustainable energy management. He is the technology lead of the award-winning La Trobe Energy Analytics Platform (LEAP), and his work spans applied AI, machine learning and building practical AI capability in industry settings. Connect with Nishan Mills on LinkedIn	Workshops
12:00-13:00	📍 Lunch/Coffee Area Lunch	Lunch
13:00-15:00	📍 301 Agent Skills: Building and Adopting AI Agents for Workplace Productivity Agent Skills: Building and Adopting AI Agents for Workplace Productivity What AI agents are, where they are useful, and how organisations can adopt them without over-automating critical decisions. Participants map real workplace use cases such as document handling, research, reporting, operations and internal knowledge workflows. Harsha Moraliyage is a senior AI architect and AI research scientist whose work spans Responsible AI, natural language processing, deep learning, cybersecurity and energy analytics, including generative AI agents for hyper-automation. He co-designed a university-wide AI literacy module giving coursework and research students core competence in AI, and has published widely on responsible AI principles and practice.	Workshops

Wednesday 2 September 2026

	Connect with Harsha Moraliyage on LinkedIn
15:00-15:30	📍 Lunch/Coffee Area Coffee Break Coffee
15:30-17:30	📍 301 Workshops Sovereign AI: Hosting and Running Local Models Sovereign AI: Hosting and Running Local Models A practical introduction to why local AI matters and how organisations can run models on infrastructure they control: open models, hardware choices, Ollama, LM Studio, model serving, privacy trade-offs, cost considerations, and when local AI is (and is not) the right choice. Prabod Rathnayaka is co-founder and CTO of Base Compute, where he works on local AI, helping organisations run models on infrastructure they control. A senior machine learning engineer and AI researcher with 6+ years of experience across natural language processing, voice AI and multimodal learning, he was previously an AI research scientist at Maincode, the Melbourne AI lab building Matilda, an Australian-built large language model, and has held roles at Replica Studios and John Snow Labs. Connect with Prabod Rathnayaka on LinkedIn
17:30-19:30	📍 301 Workshops AI Policy and Risks: How to Write an AI Policy for Your Organisation AI Policy and Risks: How to Write an AI Policy for Your Organisation A parallel governance-focused session for leaders, managers and teams responsible for safe AI adoption: acceptable use, data handling, approval processes, human oversight, tool governance, risk classification and staff training. Zafar Issadeen is an AI and construction technology consultant who advises organisations on adopting AI and digital tools responsibly across construction and infrastructure. He brings deep construction-tech experience — over 2,000 projects delivered across 12 years — including senior technology leadership at Victoria's Level Crossing Removal Project and the delivery of collaboration platforms for the \$10.9B Metro Tunnel, with earlier roles in project delivery at InEight and ICT advisory at KPMG covering security, privacy and continuity. He is a Certified Information Systems Auditor (CISA) and Certified Information Systems Security Professional (CISSP). Connect with Zafar Issadeen on LinkedIn

Wednesday 2 September 2026

📍 201

Workshops

Responsible AI in Practice: Deploying Generative and Applied AI Systems for Real-World Impact

The rapid adoption of Artificial Intelligence, particularly Generative AI, has intensified the need for practical, deployable approaches to Responsible AI. While principles such as fairness, transparency, safety, and accountability are widely discussed, there remains a gap between theory and real-world implementation.

This tutorial presents a structured, practice-oriented framework for designing, evaluating, and deploying AI systems responsibly across diverse domains. Drawing on interdisciplinary research and real-world deployments from the Multidisciplinary AI Research Centre (MARC), University of Peradeniya, the session focuses on how Responsible AI can be operationalized in education, healthcare, agriculture, climate, and public-sector contexts.

Workshop Agenda (2 Hours)**1. Setting the Context (15 min)** [Prof Vijitha Herath]

- From AI models to AI systems and agents
- Why Responsible AI is now a deployment problem
- Key risks: bias, hallucination, misuse, overreliance

2. Core Pillars of Responsible AI (20 min) [Prof [Roshan Ragle](#)]

- Fairness, transparency, safety, privacy, accountability
- Trade-offs between performance and responsibility
- Regulatory and institutional perspectives

3. Responsible AI Lifecycle for Deployment (30 min) [Prof [Roshan Godaliyadda](#)]

- Problem framing and use-case selection
- Data, model, and system design choices
- Evaluation before deployment
- Monitoring, feedback loops, and governance

4. Case Studies from MARC (30 min) [3 RAs from MARC]

- AI in healthcare
- AI in agriculture and climate
- AI in education
- Lessons from real deployments

6. Institutional Readiness and Scaling (25 min) [Dr [Asitha Bandaranayake](#)]

- Skills, governance, and policy requirements
- Scaling responsible AI adoption
- Assess a sample AI system using a Responsible AI checklist (Interactive Evaluation Exercise)

Thursday 3 September 2026

09:00-09:30	📍 Lunch/Coffee Area Reception Desk + Breakfast Coffee
09:30-12:00	📍 101 Papers AI Agents and Systems 1 09:30-09:45 IRAI26-000178 - Who Is Responsible When AI Decides? A Pre-Deployment Accountability Allocation Framework for Agentic Generative AI Systems <i>Shiv Banjade, Sangita Pokhrel, Sreekanth Manchala, Swathi Ganesan</i> 09:45-10:00 IRAI26-000025 - Autonomous Creation of Digital Twins for Hydrogen Fuel Cell Systems using Agentic AI <i>Pariya Hazratianlotfabadi, Hamid Abdi, Saman A. Gorji, Joseph Varga, Michael Pereira</i> 10:00-10:15 IRAI26-000215 - An Agentic Generative AI Framework for Industrial Predictive Maintenance: Physics-Aware Anomaly Detection with Multimodal Retrieval-Augmented Generation <i>Minindu Abeywardena, Imansha Dilshan, Supun Tharaka, Nandun Dananjaya, Danindu Ransika</i> 10:15-10:30 IRAI26-000197 - ASAB: Adversarial Sampling At Code Boundary <i>Yassine Akhiat, Yousif M. Elsheikh, Henry Späth, Bjoern Annighoefer, Zamira Daw</i> 10:30-10:45 IRAI26-000099 - Neuron Incidence Redistribution for Fairness in Medical Image Classification <i>Abin Shoby, Lyle John Palmer, Nikhil Cherian Kurian</i> 10:45-11:00 IRAI26-000127 - Edge-Cloud Collaboration in Online Class-Incremental Continual Learning <i>Kevin Shen, Yao Du, Cyril Leung</i> 11:00-11:15 IRAI26-000308 - Comparing Retrieval-Augmented Generation Configurations under Noisy Retrieval: A Taxonomic Review of Hallucination Behaviour <i>Anmita Das, Shoumya Chowdhury, Sushanta Paul</i> 11:15-11:30 IRAI26-000353 - STaR-FC: State-Transition and Reconfiguration Functional Connectivity Framework for Parkinson's Disease <i>Tayyaba Nosheen, Sohail khalid, Mateusz Janecki, Justyna Romanek, Marta Chmura, Jacek Jurkoj rPiotr Wodarski, Aamir Saeed Malik</i> 11:30-11:45 IRAI26-000391 - How to design algorithms that generate counterfactuals in Explainable AI <i>Vinura Galwaduge, Jagath Samarabandu</i> 11:45-12:00 IRAI26-000051 - Generative Demand Modeling for Trustworthy and Robust Mobile Charging Fleets <i>Ubaid Qureshi, Neha Tak</i>
	📍 201 Papers AI Applications in Practice 09:30-09:45 IRAI26-000015 - Who is in control? The need for explainable AI <i>Svetlana Bialkova</i> 09:45-10:00 IRAI26-000096 - Regulating Civilian AI Technology with Dual Use Military Applications <i>Francis Borok</i> 10:00-10:15 IRAI26-000129 - Human vs AI comparative acoustic reference-based evaluation of Naturalness and Expressiveness in the Kazakh speech <i>Aidana Karibayeva, Vladislav Karyukin, Balzhan Abduali</i> 10:15-10:30 IRAI26-000103 - Responsible Deepfake Video Detection: An Attention based Multi-Model Feature Fusion Framework Supporting Transparency <i>Sujal Upadhyay, Deepani Guruge, Nisha Kumari</i> 10:30-10:45 IRAI26-000290 - CARE-AI: A Co-Created Pedagogical

Thursday 3 September 2026

Framework for Responsible AI Learning Experiences for Young Children *Atm Shafiul Alam, Rokaiya Sultana***10:45-11:00 IRAI26-000076 - Workout Injury Prediction from Personal Data Using ML Models** *John Ayode, Abhishek Basnet, Ashraf Uddin***11:00-11:15 IRAI26-000225 - Label-Free Detection of Incipient Electro-Mechanical Faults via Cross-Modal Attention and Latent Diffusion** *Nunsi Shiaki, Oluwaferanmi Oladepo, Adejire Adegite***11:15-11:30 IRAI26-000293 - OpenClaw-Driven Dynamic Agent Orchestration for Deep-Tier Supply Chain Disruption Intelligence** *Jeremy Varghese, Rishabh Pandya, Lakshitha Gunasekara***11:30-11:45 IRAI26-000175 - Fitness to Clinical Prescription: Fine-Tuning Vision Language Models for Therapeutic Exercises** *Gayanuka Amarasuriya, Malitha Prabhashana, Chathura Jayasinghe, Nidhan Samarasinghe, Velmanickam Logeeshan, Navien Nair, Chathuranga Ranasinghe*

📍 301

Papers.

AI Models and Algorithms**09:30-09:45 IRAI26-000333 - Forecast-Aware Reinforcement Learning for Autoscaling in Containerized Systems** *Gangani Ariyaratne, Isuru Ariyaratne, Lakshitha Gunasekara***09:45-10:00 IRAI26-000111 - Fast Retrieval, Unequal Exposure: Evaluating Efficiency and Exposure Bias in Vector-Based Recommender Systems** *Enock Ayiku, Jain Navkar, Kenneth Fletcher, Tarikul Islam Papon***10:00-10:15 IRAI26-000062 - LiveMem: Runtime Continuous Memory Management for Resilient Agentic Large Language Model Systems** *Nitin Lohana, Mahendra Rajwade, Bala Vikram Tadikonda, Rahul Padwani***10:15-10:30 IRAI26-000203 - A Windows-Native Contextual AI Assistant for Perceived Reduction of Context Switching in Novice Programming Workflows** *Thanoban Kamalendran, Udula Harith***10:30-10:45 IRAI26-000079 - Generative AI Adoption in Lower-Middle-Income Countries: Insights from Bhutan** *Phub Namgay, Pema Wangdi, Tandin Penjor, Tashi Wangchuk***10:45-11:00 IRAI26-000130 - Adaptive and Explainable Hybrid Framework for Streaming Financial Fraud Detection Under Concept Drift** *Stephen Iheagwara, Henry Mafua, Victoria Nyam, Fatimah Adamu-Fika***11:00-11:15 IRAI26-000278 - Quantum-Enhanced Variational Classifiers for Ethical Bias Detection in Multimodal AI Models** *Kavach Shah, Nikunj Doshi, Shikhar Patel***11:15-11:30 IRAI26-000124 - FAIRGate: Agentic Runtime Policy Selection and Verified Repair for Fairer Recommendation Delivery** *Navkar Uttamchand Jain, Enock Ayiku, Kenneth Fletcher***11:30-11:45 IRAI26-000098 - A Graph-Based Semi-Supervised Learning Framework for Illicit Bitcoin Transaction Detection** *Ashim Maharjan, Md Ashraf Uddin, Nam Hoai Chu, Reza Rafeh***11:45-12:00 IRAI26-000091 - A Cascaded Edge-Cloud Architecture for Automated Diabetic Retinopathy Screening** *Nishi Miteshkumar Doshi, Shrey Shah*

📍 401

Papers.

AI Safety and Security**09:30-09:45 IRAI26-000125 - Tool or Agent? The Discursive**

Thursday 3 September 2026

	Construction of AI Agency <i>Lev Sukherman, Yetunde Longe-Folajimi, Marina Konkol</i> 09:45-10:00 IRAI26-000095 - Fairness in AI-Assisted Public-Interest Decision-Making Across Employment, Credit, and Criminal Justice <i>Enock Ayiku, Sonia Arthur, Kenneth Fletcher</i> 10:00-10:15 IRAI26-000005 - Designed Affective Stewardship: Toward Life-Centered Emotional Interaction in Conversational AI <i>Eunna Lee</i> 10:15-10:30 IRAI26-000256 - Bidirectional Mamba Framework for Anomaly Detection in IIoT Systems <i>Noureen El-Shenawi, Basmala Taha, Eman Atef, Ahmed Zaky</i> 10:30-10:45 IRAI26-000167 - Test-Time Compute and Semantic Safety Filtering for Industrial Fault Attribution <i>Mohamed Alwathiq Ali Abdallah Babiker, Moaz Jalal</i> 10:45-11:00 IRAI26-000170 - Hybrid ML-LLM Framework for Partial Discharge-based Maintenance of XLPE Underground Power Cables <i>Alamera Nouran Alquennah, Mohammad AlShaikh Saleh, Sayed Mohammad Kameli, Mohammed Al-Hajri, Ali Ghrayeb, Haitham Abu-Rub, Sunil Khatri, Marek Olesz</i> 11:00-11:15 IRAI26-000184 - Multimodal Generative AI Model for Explainable PCB Defect Inspection and Automated Quality Documentation <i>ADRIAN PINADO COTRINA, D'Alessandro del Piero Sarmiento Ayala, Angello Falla Dapello, Laleska Arroyo Aquino, Mauricio Sandoval Arrieta</i> 11:15-11:30 IRAI26-000236 - Operational Regime-Aware Degradation Simulation for Mission Planning Support <i>Digvijay Rajoria, Chathurika Wickramasinghe</i> 11:30-11:45 IRAI26-000166 - Counterfactual Generation for Representation Learning in Industrial Time-Series Prognostics <i>Saverio Napolitano, Nicola Ricciardi</i> 11:45-12:00 IRAI26-000012 - Responsible AI in Automotive Diagnostics: Human-in-the-Loop Governance for AI-Driven Fault Detection <i>Ashish Rai, Jordie McAuley, Amardeep Singh, Sahar Kokaly</i>
12:00-13:00	📍 Lunch/Coffee Area Lunch Lunch
13:00-15:00	📍 401 AI Agents and Systems 2 13:00-13:15 IRAI26-000108 - Responsible Agentic AI at the Protocol Layer: An MCP Governance Gateway for Evidence Capture, Oversight, and Auditability <i>Heidrun Muhle, Andreas Stockl</i> 13:15-13:30 IRAI26-000311 - Closing the Consumer Governance Gap: Framework and Empirical Validation for AI Agent Safety <i>Abhishek Kumar Singh</i> 13:30-13:45 IRAI26-000378 - A Generative AI Multi-Agent Framework for Autonomous Virtual Net Metering in Decentralized Microgrids <i>Mercy Otuko, Federation Chitoma, Mercy Runo, IAN CHEGE, Nickson Kiprotich</i> 13:45-14:00 IRAI26-000065 - TrustLens: A Four-Pillar Testing Framework for AI-Integrated SaaS Applications <i>Vinil Pasupuleti, Koushik Kumar Ganeeb, Nagesh Gulkotwar</i> 14:00-14:15 IRAI26-000212 - EnhancedDALNet: Physics-Informed Diffusion Model for Multi-Source Power Generation Forecasting <i>Namina Wijetunga</i> Papers.

Thursday 3 September 2026

14:15-14:30 **IRAI26-000243 - PE-GPT 2.0: A Trustworthy Multi-Agent Design Framework for Power Converters** *Qiu Shi, FANFAN LIN, Xinze Li*
 14:30-14:45 **IRAI26-000121 - SAFE-LLM: Scenario-Aware Counterfactual Fairness Verification for Large Language Model Decision Support** *Enock Ayiku, Navkar Jain, Kenneth Fletcher*
 14:45-15:00 **IRAI26-000144 - ReSceneScribe: Static-Dynamic 3D Accident Reconstruction from Vehicle-Mounted Cameras and Multi-Agent Simulation for Evidence-Grounded Investigation** *WOOSUP LEE, YONGMIN KIM*

📍 201

Papers

Human-AI Systems and Practice

13:00-13:15 **IRAI26-000227 - AviAgent: An Agentic System for ICAO-Aligned Aviation Safety Report Classification** *NASIRU Iliya, Manyara Baraka, Delyon Ronnie, John Bosco Thuo, Joao Barros*
 13:15-13:30 **IRAI26-000092 - The Feedback Lens: A Trustworthy AI Feedback System for Higher Education** *Haoyu Li, Gangsen Zhou, Melvin Saji, Charith Jayasekara, Lakshitha Arachchige*
 13:30-13:45 **IRAI26-000078 - When Accuracy Is Not Enough: Rationale Extraction for Vulnerable Customer Protection in Electricity Disconnection Decisions** *Fil Prager, Charith Jayasekara, Sean Paterson*
 13:45-14:00 **IRAI26-000031 - MedAdapt-SAM: Hybrid Prompt-Driven and Uncertainty-Aware Adaptation of Foundation Models for Brain Tumor Segmentation** *Shaswat Pradhan, Nikhil Kumar, Sunil Mohapatra, Sujata Chakaravarty, Chinmayee Dora*
 14:00-14:15 **IRAI26-000018 - Bridging Navigation and Manipulation via Whole-Body Motion for Mobile Manipulators** *I-En Chiu, Weicheng Wang, Rongshun Chen*
 14:15-14:30 **IRAI26-000024 - FAIRbias: A Fine-Tuned DistilBERT for Binary Bias Detection** *Matthew Chan, Arun Kumar Rajasekaran*
 14:30-14:45 **IRAI26-000327 - Human-AI Workforce Collaboration: The Junior API & Parser Programmers, the “fitting-the-pipes” Analogy Between the Residential & Industrial Plumbers (AI in Using or Making?)** *Yeeshtdevisingh Hosanee*

📍 101

Papers

Multimodal AI and Industrial Systems

13:00-13:15 **IRAI26-000153 - Data-Driven Cryptographic Agility with Bounded\AI Decision Models for Industrial IoT Systems** *Abdullah Aydeger, Engin Zeydan*
 13:15-13:30 **IRAI26-000128 - Structuring Individual Model for Force Estimation from Surface EMG Using Responsible Artificial Intelligence** *Daiki Sodenaga, Daswin de Silva, Seiichiro Katsura*
 13:30-13:45 **IRAI26-000282 - Natural-language based system maximizing explainability for 3D space using vision-language model** *Taiga Sekida, Shinnosuke Sato, Ryusei Sugano, Ryoichi Shinkuma, Gabriele Trovato*
 13:45-14:00 **IRAI26-000043 - Audio Source Separation and Automatic Music Transcription on Specific Instruments** *Saneru Akarawita, Manjusri Wickramasinghe, Roshan Abeyweera*
 14:00-14:15 **IRAI26-000072 - Securing Autonomous Rail Operations: AI and Cybersecurity Against Insider Threats** *ALI SHIKAI SHIKAI, Tim Nedyalkov*

Thursday 3 September 2026

- 14:15-14:30 **IRAI26-000073 - HiRAS: LLM-Assisted Claim Validation for Cost-Sensitive Retrieval Adaptation in Engineering Operations** *Alexander Laurence, Yoichi Kawachiya, Keiko Tanigawa, Yuuya Isoda, Hiroaki Kounoura*
- 14:30-14:45 **IRAI26-000309 - Heterogeneous Threshold Sensing with Hysteresis-Bounded Priority Switching for Autonomous Energy Management in Smart Microgrid** *Asit Kumar Sahoo, Ranjeeta Patel, Subodh Kumar Mohanty, Chathurika S. Wickramasinghe Brahmana, Daswin De Silva*
- 14:45-15:00 **IRAI26-000268 - ClaimGraph: A Multi-Detector Forensic Pipeline with LLM Audit Reasoning for Automotive Warranty Claim-Pack Fraud Detection** *Puneet Samnani, Vibhu Goenka, Shaswat Shrivastav, Vidyasree Vankam, Dipesh Joshi, Harsha Moraliyage*

📍 301

Papers.

Physical AI and Robotics

- 13:00-13:15 **IRAI26-000176 - Physics-Regularized Generative Framework for Real-Time Singularity-Aware Robot Motion Planning** *Anuththara Bombuwala Disawage, Muditha Pahuruthota Acharige, Mayumi Kasthuriarachchige Don, Dinusha Ukwaththage, Sachini Paththini Kapuge*
- 13:15-13:30 **IRAI26-000162 - T-Script: Attention based learning for Reconnaissance of Wireless Sensor Networks** *JAYANTHI GANAPATHY, Alwyn Sharon, Yogeshwaran CS, Mritula Shankar, Varsha Venkatesan*
- 13:30-13:45 **IRAI26-000069 - Autonomous API Sandbox Generation Using Agentic AI** *Swapneswar Sundar Ray*
- 13:45-14:00 **IRAI26-000315 - PriceReason: A ReAct-Based LLM Agent for Explainable Dynamic Pricing in Retail** *Mitchelle Ashley Creado*
- 14:00-14:15 **IRAI26-000330 - Positive Backward Transfer Under Family-Grouped Task Sequences in Replay-Based Continual Learning: A Confirmatory Study on an Abstract-Reasoning Benchmark** *Lava Kumar Polu, Maria Vaida, Aditya Syal*
- 14:15-14:30 **IRAI26-000381 - cWGAN-GP-Based Minority-Class Augmentation for Edge-Efficient Industrial Defect Classification** *Salah Eddine Mamouni, Anass Mamouni*
- 14:30-14:45 **IRAI26-000312 - Soft Failures in Enterprise AI: A Framework for Organizational Risk in LLM-Assisted Workflows** *Swapneswar Sundar Ray*
- 14:45-15:00 **IRAI26-000247 - TrustEx: Authority-Bounded Generative Reasoning over Trust-Stratified Heterogeneous Signals for Real-Time Industrial Exception Management** *Sandar Fernando, Harsha Moraliyage*

15:00-15:30

📍 Lunch/Coffee Area

Coffee

Coffee break

15:30-16:30

📍 101

Plenary

Keynote 1 - ARM and the Future of AI Compute
Keynote 1 - ARM and the Future of AI Compute
Speaker - Robert Iannello and Matt Cossins

Thursday 3 September 2026



Robert Iannello
Senior Manager,
Academic and Emerging
Ecosystems, Arm



Matt Cossins
Developer Ecosystem
Manager, AI & Developer
Platforms, Arm

AI inference is rapidly expanding from the cloud to the edge, enabling increasingly intelligent experiences across mobile, automotive, industrial, IoT, and embedded systems. This session explores Arm's AI at the Edge strategy and how technologies such as ExecuTorch are helping translate that strategy into practical on-device AI deployment.

The session begins with a 30-minute in-person presentation from Robert Iannello on Arm's AI at the Edge strategy, exploring the evolution of AI from cloud-based computing towards increasingly distributed and on-device experiences. It provides an overview of Arm's approach to enabling AI across a broad range of devices and markets, highlighting the importance of efficient, scalable computing and a strong software ecosystem. It also considers how advances across hardware and software are enabling developers to bring increasingly capable AI experiences to the edge.

Building on this strategic context, a technical presentation by Matt Cossins introduces ExecuTorch and its role in extending the PyTorch ecosystem to edge and embedded devices. The presentation follows the journey from PyTorch model development towards efficient on-device execution, providing an overview of model preparation, ahead-of-time compilation and optimization, runtime deployment, and backend delegation. Together, the two presentations connect the broader evolution of AI at the edge with the practical software workflows developers can use to deploy AI applications across Arm-based devices.

16:30-17:30

📍 101

Plenary

AI Founders Panel: Building on Shifting Ground – What Survives the Next Model Release


AI Founders Panel: Building on Shifting Ground – What Survives the Next Model Release

Perspectives from founders building on rapidly evolving AI infrastructure

Every six months, a new frontier model release erases someone's product roadmap and creates someone else's company. This panel brings together founders at different stages of that gamble. They'll talk candidly about where defensibility actually comes from when the underlying models are commodity infrastructure: proprietary data, workflow depth, distribution, regulatory access, or something less obvious. Expect discussion of the unglamorous parts of evaluation, reliability, enterprise procurement, unit economics that only work at scale, alongside the strategic bets that don't

Thursday 3 September 2026

show up in a pitch deck.



Sam Saltis

CEO, Core dna

Sam Saltis is an Australian technology entrepreneur, executive, and the Founder and CEO of Core dna, a SaaS-based digital experience and e-commerce platform. With over 25 years of experience in the web technology and digital strategy sectors, he has built a reputation as a thought leader in the Australian tech sector.

Dr. Jing Mu

Co-Founder, Virtetic

Dr Jing Mu is an Australian-based biomedical engineer, academic, and entrepreneur who serves as the Co-Founder and Chief Research Officer of Virtetic, a digital health and medtech company.

Dr. Ravini Savindya

Co-Founder, QuantCare

Dr Ravini Savindya is a psychologist and technology entrepreneur. She is the Founder and CEO of QuantCare, a Melbourne-based digital health analytics startup that helps medical practices deliver value-based care and improve patient outcomes.

Dr. Lukas Wesemann

CEO and Co-Founder, Base Compute

Dr. Wesemann is the Co-Founder and CEO of Base Compute, a startup dedicated to building the infrastructure required to run artificial general intelligence (AGI) locally on-device. Dr Wesemann transitions fluidly between hardware and software, leveraging a background in nanotechnology to accelerate computer vision and AI performance.

Kanchana Wickremasinghe

Vice President & Field Chief Technology Officer, WSO2

Thursday 3 September 2026

	Kanchana Wickremasinghe is the Vice President & Field Chief Technology Officer at WSO2, a global leader in digital transformation, API management, and integration technology. Known for his advocacy of open-source ecosystems, he specialises in Kubernetes architectures, internal developer platforms (IDPs), and cloud-native engineering. Dr. Harsha Moraliyage AI Architect, La Trobe University Dr. Harsha Moraliyage is a senior AI architect and AI research scientist whose work spans Responsible AI, natural language processing, deep learning, cybersecurity and energy analytics, including generative AI agents for hyper-automation.
17:30-19:00	📍 101 SYP Event and Awards Ceremony Plenary
19:00-21:00	📍 101 AI Trivia Night with Pizza Dinner

Friday 4 September 2026

09:00-09:30	📍 Lunch/Coffee Area Reception Desk + Breakfast Coffee
09:30-11:00	📍 101 Opening Ceremony + Keynotes IRAI 2026 Keynote Speakers: Satya Tammareddy, Head of GTM ANZ, OpenAI Craig Scroggie, CEO and Managing Director, NEXTDC Professor Theo Farrell, Vice-Chancellor, La Trobe University  Satya Tammareddy Head of GTM, ANZ OpenAI  Craig Scroggie CEO and Managing Director NEXTDC  Professor Theo Farrell Vice-Chancellor La Trobe University Satya Tammareddy, Head of GTM ANZ, OpenAI Satya partners with organisations of all sizes - from startups to large

Friday 4 September 2026

enterprises - to drive AI adoption at scale and deliver measurable business impact. Prior to OpenAI, Satya worked at Stripe building and scaling high growth businesses across APAC.

OpenAI is the world's leading frontier AI lab. Its core mission is to ensure that artificial general intelligence benefits all of humanity. <https://openai.com/>

Craig Scroggie, CEO and Managing Director, NEXTDC

Craig is the CEO and Managing Director of NEXTDC (ASX:NXT). Craig has more than 30 years' experience in the deep tech sector, having held senior positions with Symantec, Veritas Software, Computer Associates, EMC Corporation and Fujitsu. Craig is also Chair of Data Centres Australia and serves on the Board of Freelancer (ASX:FLN).

NEXTDC is Australia's leading AI infrastructure provider, an ASX-100 company that designs, builds, and operates premium data centres for critical digital infrastructure. <https://www.nextdc.com/>

Professor Theo Farrell, Vice-Chancellor, La Trobe University

Professor Theo Farrell is Vice-Chancellor and President of La Trobe University. A war studies scholar by background, Professor Farrell is a leading expert on strategy and organisational innovation. He is an influential voice on how universities can adapt to technological change and deliver greater value to students and society.

La Trobe University is Australia's AI-first university, serving 40,000 students across seven metropolitan and regional campuses. La Trobe is ambitiously transforming research, education, industry engagement and student support through a responsible AI-first adoption strategy. <https://www.latrobe.edu.au/>

11:00-11:15

📍 Lunch/Coffee Area

Coffee

Coffee Break

11:15-12:30

📍 101

Papers

AI Literacy and Education

Chairs: Alexia Maddox, Shalinka Jayatilleke

11:15-11:30 IRAI26-000093 - A Community Co-Design Model for AI Literacy and Technology-Facilitated Abuse Prevention in CALD Communities *Alexia Maddox, Haroon Sayed, Druma Datey, Sarah Barnbrook, Jenine Beekhuyzen, Supriya Singh*

11:30-11:45 IRAI26-000306 - When the Regulator Cannot Read the Model: Technical Governance Capacity Gaps and AI Safety in Developing Countries *Shoumya Chowdhury, Anmita Das, Sushanta Paul*

11:45-12:00 IRAI26-000016 - Structured Integration of Generative AI in Authentic Learning and Assessment with Cognitive Perspectives *Charith Jayasekara, Lahiru Akuratiyage, Carlo Kopp, Chetan Arora, Vincent Lee, Chamila Gunasekara*

12:00-12:15 IRAI26-000168 - Towards Responsible AI in Education: Explainable Adaptive Learning through Active Inference *Pengyi*

Friday 4 September 2026

Wang, Atm Shafiul Alam

12:15-12:30 IRAI26-000371 - **Operationalising Whole of Institution GenAI Governance: A Public Evidence Diagnostic for Higher Education** Shalinka Jayatilleke, Aneela Yasmin, Rashikala Weerawarne, Akhtar Jalbani, Akrtati Jain

📍 201

Papers

AI Safety and Assurance

11:15-11:30 IRAI26-000063 - **Seatbelt: An Open-Source Python Library for Responsible AI Auditing of Generative AI Models** Michelle Lee

11:30-11:45 IRAI26-000086 - **LLM-ASG: Policy-Driven Controls to Prevent Prompt Injection and Data Leakage** Tamirat Abegaz, Orrin Adotevi

11:45-12:00 IRAI26-000070 - **Mitigating Hallucination and Label Inflation in LLM Classification of Mine Rehabilitation Technical Documents via Prompt and Class-Scheme Refinement** Josh Koh, Charlie Taylor, Gabriel Maicas

12:00-12:15 IRAI26-000332 - **Artificially Influenced: The Effect of AI Agent Consensus and Anthropomorphism on Human Conformity** Constance Holmes, Thom Griffith, Phillip Morgan

12:15-12:30 IRAI26-000106 - **Compliance as Architecture: A Regulatory Compliance Traceability Matrix and Overhead Classification for Runtime AI Governance** MARIA LUZ MADARIAGA

📍 301

Papers

Industrial AI and Intelligent Operations

11:15-11:30 IRAI26-000206 - **LLM-Based Rapid Generation Method for Power Distribution Networks Service Restoration** Haidian Li, Xupeng Liu, Yuan Yao, Junwei Yang, Zhiqun Zhang, Bing Sun, Hao Yu, Haoran Ji, Di Wang, Guanyu Song, Jinli Zhao, Peng Li

11:30-11:45 IRAI26-000100 - **Graph-Enhanced Multi-Scale Transformer with Adaptive Routing for Silicon Content Prediction in Blast Furnace Ironmaking** Zhanghui Deng, Qiang Liu, Yishuai Liu

11:45-12:00 IRAI26-000241 - **A Generative AI Integrated Multimodal Framework for Low-Latency Multi Camera Person Re-Identification** Leon Fernando, Charitha Dombawala, Pramod hettigoda, Ishara Neranjana, Vanodhya Warnasooriya, Rashmika Nawaratne

12:00-12:15 IRAI26-000188 - **A Framework for Guiding LLMs to Solve Constrained Combinatorial Optimization Problems: A CVRP Case Study** Shuang Liu, Tharindu Bandaragoda

12:15-12:30 IRAI26-000179 - **MARVEL-RAG: Product-Grounded Multimodal Retrieval-Augmented Generation for Automated Social Media Marketing Content in E-Commerce** Sangita Pokhrel, Nalinda Somasiri

📍 401

Papers

Physical AI and Robotics

11:15-11:30 IRAI26-000019 - **LLM-Based Target Object Grasping and Packing Strategy Planning in Cluttered Environments** Yanxiang Fang, Weicheng Wang, Rongshun Chen

11:30-11:45 IRAI26-000216 - **Perpetual Cognition for Real-Time Collaboration: Graph-Based Reasoning in Human-Machine**

Friday 4 September 2026

	Teamwork <i>Thien Huynh, My Lai, Sudheera Withanage</i> 11:45-12:00 IRAI26-000017 - A Vision-Based Framework for Cross-Device Topological Localization and Navigation <i>Jing Chen, Weicheng Wang, Rongshun Chen</i> 12:00-12:15 IRAI26-000020 - Structural Representations for Cross-Attack Generalization in AI Agent Threat Detection <i>Vignesh Iyer</i> 12:15-12:30 IRAI26-000352 - A Three-Layer Institutional Complementarity Framework for AI Capability: Evidence from China, the EU, the UK, and the US <i>Xi He, Tien-Der Han, Anna Grosman</i>
12:30-13:30	📍 Lunch/Coffee Area Lunch Lunch
13:30-14:30	📍 101 Keynote Keynote by Prof. Saman Halgamuge.  Professor Saman Halgamuge The University of Melbourne Large Language Models (LLMs) have transformed the field of artificial intelligence. Their extension to computer vision has led to the emergence of Vision-Language Models (VLMs), paving the way for the next generation of Physical AI through Language-Action Models (LAMs) and Vision-Language-Action (VLA) models. These models integrate visual and linguistic information to generate concrete actions for autonomous agents and machines including robotic systems. VLA models are fundamentally changing robotics. Rather than executing pre-programmed sequences of operations in structured environments, VLA-enabled robots are expected to interpret instructions, reason using external knowledge, adapt to unfamiliar situations, and perform tasks in previously unseen environments. Concurrent advances in world models, agentic AI, reinforcement learning, and foundation models are accelerating the development of intelligent robots capable of autonomous action. The assumption that VLAs can fully represent the physical world using vision, language, and sensor inputs encoded as tokens is, at present, optimistic. Although these multimodal inputs can be fused through cross-attention mechanisms and transformer architectures can generate action sequences, there remain open research challenges.

Friday 4 September 2026

This talk will examine the current state of Physical AI, the scientific and engineering challenges that remain, and the opportunities for deploying intelligent machines in complex, dynamic, and unstructured real-world environments.

One promising avenue we pursue in our research group is the development of richer world models that integrate scientific knowledge with data-driven learning. Physics-Informed Neural Networks (PINNs) incorporate physical laws, expressed as differential equations, into machine learning models. However, a comprehensive world model should extend beyond physics and training data and the corpora of general text it was trained on (if the backbone used is an LLM). Knowledge is also encoded in subject-specific textual descriptions, knowledge graphs and other structured representations. To build more robust and generalisable world models, we should exploit all available forms of prior knowledge. We have developed a new framework that enriches parameter discovery of model equations by integrating auxiliary knowledge sources alongside observational data [1, 2].

A second challenge we pursue is the integration of continuous time-series data into VLA models. Industrial robots operate in environments rich in continuous sensor signals, including force, vibration, current, temperature, and motion measurements. However, most LLMs and VLMs are pretrained on discrete modalities such as text and images, making their adaptation to continuous temporal signals non-trivial. To address this gap, we propose a temporal modulation framework that fine-tunes LLMs and VLMs to effectively process time-series inputs alongside images and text, enabling VLA models to capture complex temporal dynamics and make more informed decisions in real-world robotic applications [3].

References:

1. Knowledge Inclusive Machine Learning for Disease Gene Prioritisation. CJ Gamage, Y Xia, R Rupasinghe, S Senevirathne, D Senanayake, ..., S. Halgamuge. bioRxiv 2026.04.29.721522, 2026.
2. Knowledge Inclusive Adaptive Physics-Informed Neural Network for Microbial Interaction Modelling. R Rupasinghe, R Vidanaarachchi, A Hevathige, S Seneviratne, S. L. Tang, S. Halgamuge. arXiv:2606.07686, 2026.
3. Rethinking Time Series Forecasting with LLMs via Nearest Neighbor Contrastive Learning. J Bogahawatte, S Seneviratne, M Perera, S Halgamuge. arXiv:2412.04806, 2026.

14:30-15:30

📍 101

Plenary


Human-AI Futures Forum

Human Skills for an AI World

This panel explores how AI is reshaping work, workplaces, jobs and skills. Stepping beyond the hype cycle of which jobs AI will replace, the panel

Friday 4 September 2026

examines which human capabilities become more valuable as AI becomes more capable. The panel brings together multidisciplinary perspectives from academia, industry and community, with panellists from AI research, AI education, AI policy and AI indigenous knowledge systems. The panel will unpack macro trends across the knowledge economy, followed by work roles, activities and tasks that are automated or human centred. Emerging indicators and insights into human-AI collaboration, AI automation and the impact on different levels of knowledge work, will then inform what human skills are critical for an AI world and how they should be identified, refined, taught and developed in educational, professional and industry settings.



Panelists:

Dr Shalinka Jayatilleke

Senior Lecturer in AI and Analytics, La Trobe University

Dr Shalinka Jayatilleke is a Senior Lecturer in AI and Analytics at La Trobe University and Course Coordinator of the Bachelor of Business Analytics. She is an educator and researcher with expertise in AI, analytics, AI literacy, software design, and digital transformation, with a strong focus on preparing students to use AI and data-driven technologies responsibly and effectively. She has co-authored more than 20 research articles, is a co-investigator on industry research grants, and currently supervises six PhD students in industry-focused applied AI topics.

Karine O'Donnell

Founder & Director, Projecting With People

Karine O'Donnell is Founder & Director at Projecting With People (PWP). PWP helps businesses structure execution, improve governance, and drive measurable success, without burning out their teams. The PWP approach blends strategy, business process engineering, practical frameworks, project delivery, change management and a people-first mindset.

Lily Ballot Jones

AI Law Researcher, KomplyAI

Lily Ballot Jones is an AI law researcher working across AI governance in academia, industry, and the public sector. At Australian startup KomplyAI, she translates global AI regulation into product insights for AI-driven compliance systems, while at Swinburne University she contributes to

Friday 4 September 2026

research on AI-resilient assessment design. Previously, Lily worked at the University of Melbourne's Centre for AI and Digital Ethics, where she led and published research on AI in workplace surveillance.

Dr Rajith Vidanaarachchi

Research Fellow, University of Melbourne

Dr Rajith Vidanaarachchi is a Research Fellow at the University of Melbourne. His research applies agent-based modelling, complex systems approaches and AI to social, spatial and health-related questions, with particular interests in fairness and ethics in AI and the relationship between AI systems and Indigenous Knowledge Systems, including issues of data sovereignty and community participation in AI development. His recent work includes contributing to research on housing insecurity and intergenerational disadvantage as part of the ARC Life Course Centre, and he has lectured and published on fairness and ethics in generative AI and large language models.

Dr Behnam Forouhandeh

Director of Academic Services, Australia, Risepoint

Dr Behnam Forouhandeh is Director of Academic Services, Australia at Risepoint. He leads the learning design and course production team behind Risepoint's Australian online programs with a focus on embedding AI and automation responsibly into learning design practice. His career in online program management spans partnerships with leading Australian and top-ranked international universities. Behnam holds a PhD from the University of Wollongong, where his research focused on marketing and the linguistic analysis of social media.

15:30-16:00

📍 Lunch/Coffee Area

Coffee

Coffee Break

16:00-17:00

📍 201

Papers

Applied AI, Agents & Education

16:00-16:15 IRAI26-000180 - Phoenix: Safe End to End GitHub Issue Resolution via Multi-Agent LLMs *kipngeno koech, Muhammad Adam, Baimam Boukar Jean Jacques, Joao Barros*

16:15-16:30 IRAI26-000010 - Equilibrium Justice in AI Systems: A Framework for Responsible Deployment *Justin Stark, ASif Gill*

16:30-16:45 IRAI26-000177 - A Generative Agentic AI Framework for Analysis and Debiasing of Vision Language Models *Jayana Gunaweera, Dileesha Rajapakse, Harsha Moraliyage, Kalana Gunaweera, Rukshan Fernando, Supipi Gunathilaka*

16:45-17:00 IRAI26-000230 - Does Specifying How Help? Model-Dependent Effects of Procedural Prompts on LLM Data Analysis Agent Stability *Taiga Yamamoto, Yukinori Sato, Kumiko Isu, Hiroko Kobori, Toshiyuki Hatta, Katsumi Okuda*

17:00-17:15 IRAI26-000269 - FIE: Generative Synthesis of Industrial Failure Scenarios for Stress-Testing Industrial AI Systems *Aahir Basu*

Friday 4 September 2026

📍 101

Papers

Industrial Applications of AI 3

16:00-16:15 **IRAI26-000158 - Proactive Flight Rerouting via GAN-Augmented Delay Prediction and Explainable LLM Advisory** *William Zhou, Huiyuan Zou*

16:15-16:30 **IRAI26-000088 - A Novel Cross-Domain RUL Prediction Method Based on Multi-Degradation Mode Perception** *Zhuo Li, Qiang Liu*

16:30-16:45 **IRAI26-000155 - QualityAgent: Multi-Agent Framework for Automated Industrial Defect Diagnosis and CAPA Generation** *Lin Luo, Amali Lakshika Mathara Arachchige*

16:45-17:00 **IRAI26-000396 - Calibrated Deep Learning for Mission-Profile-Based Reliability Estimation of Modular DAB EV Chargers** *Virginia Negri, Alessandro Mingotti, Roberto Tinarelli, Lorenzo Peretto, Riccardo Mandrioli*

17:00-17:15 **IRAI26-000395 - Physics-Guided cVAE Augmentation for Robust Degradation Diagnosis in Three-Phase Inverters** *Virginia Negri, Francesco Bazzani, Alessandro Mingotti, Riccardo Mandrioli*

📍 401

Workshops

Panel - Can AI help family violence survivors achieve 'safety'?**Can AI help family violence-survivors achieve 'safety'?**

Promise, peril, and practice at the intersection of AI and family violence response

This session addresses the question: Can AI help family violence-survivors achieve safety? We explore multiple dimensions of safety from the perspectives of family violence service providers, designers and technologists.



Prof. Supriya Singh

Chair, IEEE IC25-008;
Adjunct Professor, La Trobe
University




Sarah Barnbrook

Founder, Away from the
Keyboard




Druma Datey

Multicultural Women's
Alliance Against Family
Violence, IndianCare




Haroon Sayed

Head of Operations,
Bakhtar Community
Organisation



Zainab Sakhizada

Head of Operations,
Bakhtar Community
Organisation



Dr. Greg Adamson

Vice President, SSIT
Technical Activities

Panelists**Prof. Supriya Singh**

Chair of IEEE IC25-008; Adjunct Professor, La Trobe University and Multicultural Women's Alliance Against Family Violence

Friday 4 September 2026

Professor Supriya Singh is an Australian sociologist, writer, and academic specializing in the sociology of money, migration, and economic abuse. An Adjunct Professor at La Trobe University and Honorary Professor at RMIT University in Melbourne, her work links private financial dynamics with broader social issues.

Sarah Barnbrook

Founder, Away from the Keyboard

Sarah Barnbrook is an Australian digital safety advocate, UN delegate, and the Founder and CEO of Away from Keyboard (AFK) Inc., a charity focused on preventing technology-facilitated harm and advancing ethical AI design.

Druma Datey

Multicultural Women's Alliance Against Family Violence; and IndianCare

Druma Datey is a social welfare manager, family violence specialist, and social worker based in Victoria, Australia. She serves as a Welfare Services Manager for IndianCare, an organization supporting the Indian diaspora, and has been recognized for her advocacy and community work in the northeastern metropolitan region of Melbourne.

Haroon Sayed

Head of Operations, Bakhtar Community Organisation

Bakhtar Community Organisation is a multicultural Australian charity providing community support, education, employment, health, refugee, and inclusion services.

Zainab Sakhizada

Head of Operations, Bakhtar Community Organisation

Bakhtar Community Organisation is a multicultural Australian charity providing community support, education, employment, health, refugee, and inclusion services.

Dr. Greg Adamson

Vice President, SSIT Technical Activities

Dr. Greg Adamson is the Chief Information Security Officer (CISO) for the Victorian Department of Transport and Planning in Australia.

📍 301

Papers

Predictive, Multidisciplinary & Applied AI

16:00-16:15 IRAI26-000202 - Set-Based Multi-View Trajectory

Modeling for Uncertainty-Aware Wound Anomaly Detection *Reina*

Friday 4 September 2026

	<i>Kobayashi</i> 16:15-16:30 IRAI26-000011 - Modelling Noisy Quantum Search as an Intelligent Dynamical System with Structural Drift in Grover Dynamics <i>Ge Zhong, Charith Jayasekara, Nisal De Silva</i> 16:30-16:45 IRAI26-000122 - Causal Prior-Guided Neural Network for Time-Evolving Systems with Latent Confounders <i>Shahenda Youssef, Julius Pfrommer, Zahra Sadat Hajseyed Nasrollah, Tobias Wurth, Jurgen Beyerer</i> 16:45-17:00 IRAI26-000064 - A Responsible Generative AI Framework for Social Insurance Benefit Eligibility Assessment <i>Reshmy Krishnan</i>
17:50-22:00	📍 NGV Garden Restaurant Conference Dinner Dinner

Saturday 5 September 2026

09:00-09:30	📍 Lunch/Coffee Area Breakfast Coffee
09:30-11:30	📍 301 Responsible Decentralised Agentic Systems (ReDAS): Operationalizing Responsibility and Accountability of AI agents SPECIAL SESSION Responsible Decentralized Agentic Systems (ReDAS): Operationalizing Responsibility and Accountability of AI agents Workshops Operationalizing responsible AI spans research and development in the field from "ethics-by-design" to "assurance-by-verification." The special session is focused on a decentralised agentic system, where autonomous agents propose, argue and make decisions in a decentralised environment, subject to global AI regulatory compliance (EU AI Act, NIST RMF) and the data protection legislation (Australian Privacy Act, EU GDPR). While formal ethical principles are well-established, there is a demand for methods rigorous operational proof. The design of the session links theoretical developments with systems engineering, legal, and sociotechnical challenges of building verifiable Responsible AI (RAI) systems. We invite researchers and practitioners from the fields of agentic AI systems, complex systems, value-aligned systems, decentralised computing, AI governance and data sovereignty to explore the following interconnected pillars: • technical robustness of RDAS, focusing on deviation and hallucination detection, and consensus fail-safes for agentic workflows; • consensus in decentralised agentic systems, with embedded argumentation mechanisms for decision-making; • algorithmic fairness in value-aligned systems, including empirical bias mitigation and "unlearning" techniques for values adjustments; • human-interpretable XAI for actionable transparency in RDAS; and

Saturday 5 September 2026

- automated compliance auditing for RDAS.

Session Schedule

- 9:30 – 10:00 Explainability for Multi-Agent Recommender System by Jairo Lefebvre and Vanina Martinez
- 10:00 – 10:30 Anthropomorphic AI agency: The user perspective by Svetlana Bialkova
- 10:30 – 11:00 Decentralized AI Systems: State-of-the-art of the integration of AI with consensus systems and transaction systems by Simeon Simoff
- 11:00 – 11:30 Round table

SIMEON J. SIMOFF is Professor of Artificial Intelligence and Information Technology at Western Sydney University, where he has accumulated almost two decades of experience in various senior leadership roles. His research contributes across the fields of artificial intelligence, data science, visual analytics, design computing and augmented reality. Much of it has been supported by \$11.3mIn of competitive research funding and reported in more than 340 publications. His current work is focused on interpretable machine learning, explainable AI, and the intersection of visualisation and AI. He is founding director of the Institute of Analytics Professionals of Australia. He serves on the ACS Technical Advisory Board (AI Committee) and as Australian Chair on the IFIP Technical Committee on Artificial Intelligence (IFIP TC12). He is Fellow of the Royal Society of New South Wales and the Institute of Engineers Australia.

JAIRO A. LEFEBRE LOBAINA is a Research Assistant at the Artificial Intelligence Research Institute (IIIA-CSIC), Spain, where he works on the NL4XAI project on self-explaining multi-agent recommender systems.

His research focuses on formal explainability of black-box classifiers, efficient model-agnostic explanation algorithms, interactive explanation dialogues, and computational trust. He is a

PhD candidate at the IIIA, holds an MSc in Computer Science and Artificial Intelligence, and has prior industry experience as a data scientist.

MARIA VANINA MARTINEZ is a tenured scientist at the Artificial Intelligence Research Institute (IIIA-CSIC) in Barcelona. From 2014 to 2022 she was an adjunct researcher at the Argentine National Research Council and a professor at Universidad Nacional del Sur and Universidad of Buenos Aires.

She holds a PhD in Computer Science from the University of Maryland, College Park. In 2018, she was selected by IEEE Intelligent Systems as one of the “AI’s 10 to Watch,” an award recognising early-career researchers with outstanding contributions in AI. She has also been

Saturday 5 September 2026

awarded the 2021 Argentine National Academy Stimulus Award (Engineering Sciences) and the 2023 Konex Award in Science and Technology (Argentina).

Her research focuses on knowledge representation and reasoning, particularly in the areas of belief revision, managing inconsistency and uncertainty, responsible/artificial intelligence ethics, argumentation, and automated reasoning.

SVETLANA BIALKOVA is Professor in Information Technologies and Systems, Digital Market Applications. She is leading the ConsuWAIS lab at Sofia University, Bulgaria, with award winning grant by VIHREN 2024 program of BNSF.

Professor Bialkova received several awards, among which JSPS fellowship, BSI RIKEN, Japan. Holding a PhD from the Radboud University, The Netherlands, and securing research funding, she has worked on numerous projects in top international labs in The Netherlands, Belgium, Germany, Denmark, Switzerland, UK, Japan. She explores the joint interaction between human, media, and technology, and how this influences the multisensory experience. Her work resulted in more than 100 publications, including highly ranked journals in the areas of Consumer behaviour, Digital Marketing, Psychology, Human-Computer Interaction. Professor Bialkova is the author of the books *Digital Consumer Psychology* and *The Rise of AI User Applications: Chatbots Integration Foundations and Trends*.

09:30-12:00

📍 401

Papers.

AI and Society

09:30-09:45 **IRAI26-000242 - AfriMaint: A Generative Intelligence Framework for Predictive Maintenance in Resource-Constrained Industrial Environments** Linet Gitonga, Brian Kasina, Adrine Kibandi, Latasha Ngatia, Sandra Langat, Vidura Sumanasena

09:45-10:00 **IRAI26-000222 - GenSafeBattery: A Responsible Generative Framework for Early Thermal Runaway Prediction in Electric Vehicle Batteries Operating in Tropical Regions** Ann Karagwa, Ian Ahereza, Crystal Zawedde, Nishan Mills

10:00-10:15 **IRAI26-000182 - Safety-Bounded Generative Action Sampling with Execution-Time Safety Enforcement in Autonomous Robotic Navigation** Tanneeru Sai Mahimanvitha, Vellanki Tanay, Ruthraj Gullapelli, Komrewar Rishieth, Prasad Hettiarachchige

10:15-10:30 **IRAI26-000148 - Compositional Class-Conditioned Synthesis of Plantar Thermograms for Diabetic Foot Classification** Ahmad AlShouli, Hussain AlQattan, Mohammad Abulhasan, SaidMohammed AlSaleh, Faisal AlDhafeeri, Khaled Chahine

10:30-10:45 **IRAI26-000302 - AgentSentinel: A Multi-Agent Framework for Autonomous Root Cause Analysis and Self-Healing in Distributed Systems** Hazaifa Tahir, Saboor Ahmed, Muhammad Taha

10:45-11:00 **IRAI26-000252 - AMM-RAG: An Adaptive Agentic Multimodal Optimization Framework for Industrial Data-Driven RAG with Dynamic Retrieval Policy Learning** Mohammad Hammad Ahmad, Urmit Mahida, Harshit Joshi

Saturday 5 September 2026

- 11:00-11:15 **IRAI26-000181 - TaproQA: An Agentic Generative AI Framework for Interpretable Educational Question Answering in Low Resource Languages** *Yohan Abhishek, Deshan Sumanathilaka, Archchana Sindhujan*
- 11:15-11:30 **IRAI26-000185 - Adaptive On-Device Small Language Models for Offline Industrial Maintenance Assistance with Generative Fault Scenario Synthesis** *Abhishek Bidare, Prajakta Dalawai*
- 11:30-11:45 **IRAI26-000204 - ADAPTOPT: Graph-Grounded LLM Constraint Learning for Responsible and Adaptive Vehicle Routing in Logistics** *Rishyanth Reddy, Eswar Vajja, Akash Reddy*
- 11:45-12:00 **IRAI26-000194 - Adversarial Aware Synthetic Failure Generation for Robust Predictive Maintenance in Industrial Control Systems** *Ninad Deo, Shradha Chavan*
- 12:00-12:15 **IRAI26-000389 - Governance-Aware Automated Deployment Pipelines for Compliance-Critical Financial Services Infrastructure** *Rambabu Budithe, Nagadhara Harini Kanakala*
- 12:15-12:30 **IRAI26-000399 - PHYSICIAN: A Deterministic Middleware Kernel for Mitigating Kinematic Hallucinations in Mobile Locomotion and Manipulation Domains** *Ayatullah Hanif Showunmi, Nafisat Ozohu Abdulsalam, Abdullahi Asimiyu*
- 12:30-12:45 **IRAI26-000255 - Neuro-Symbolic Multi-Agent Frameworks for Causal Root Cause Analysis** *Tanishq Singh Sisodiya, Subhanan Chatterjee, Samarth Agrawal, Divyansh Motwani, Sai Prashanth Mallellu, Rajanikanth Aluvalu*

📍 201

Papers.

AI Models and Algorithms

- 09:30-09:45 **IRAI26-000138 - Towards Explainable Brain Tumor Analysis: A Differentiable Biomarker-Integrated Multitask Learning Approach** *Kuk Hee, Zeeshan Abbas, Muhammad Afzal, Yeji Han, Seung Won Lee*
- 09:45-10:00 **IRAI26-000249 - Stage-Aware Multi-Source RAG for LLM-Based Synthesizable Verilog Generation** *Artur ZIAZETDINOV, Vladimir Zunin, Pavel Pashentsev*
- 10:00-10:15 **IRAI26-000190 - Breaking the Filter Bubble: A Semantic Pareto-DQN Framework for Multi-Objective Recommendation** *Claudio Lucio do Val Lopes, Lucca Machado, André Brandão*
- 10:15-10:30 **IRAI26-000248 - Fair-RLVR: Reducing Stereotype Bias in LLM Reasoning via Verifiable Benchmark Rewards** *Salo E S, Arjun G Ravi, Devanand A, Rini Susan V S*
- 10:30-10:45 **IRAI26-000300 - GenEdge: Generative AI Framework for Automated Bias Detection and Debiasing in Machine Learning Systems with Multi-Agent Intelligence** *Muhammad Asad, Muhammad Salman Khalid Awan, Muhammad Hamza, Maria Tariq*
- 10:45-11:00 **IRAI26-000276 - Responsible Generative AI for Mental Health: A Safety-Constrained Human-in-the-Loop Multimodal Framework for Adaptive Creative Arts Therapy** *Kitty Liu, Shaddy Saba*
- 11:00-11:15 **IRAI26-000146 - Trust Calibration in Generative AI (GAI) Systems: Confidence and Explainability as Drivers of Over-Reliance in LLMs** *ZAIRA HASSAN Amur, Yew Kwang Hooi, Gul Muhammad, Irum Naz, Hina bhanbhro, Kamran Dahri*
- 11:15-11:30 **IRAI26-000151 - A Benchmark for Cross-Border Allergen Screening of U.S. Packaged Food Labels under Japanese Allergen Labeling Rules** *Shuhei J. Yamazaki, Takuya Maekawa*

Saturday 5 September 2026

- 11:30-11:45 **IRAI26-000156 - Benchmarking Open-Weight Foundation Models for Global AI Technical Governance** *Jason Hung*
- 11:45-12:00 **IRAI26-000246 - DatasetBench: Benchmarking Pretraining Dataset for Small Language Models under Limited Compute** *Arjun G Ravi, Salo E S, Devanand A, Rini Susan V S*
- 12:00-12:15 **IRAI26-000295 - SAFEagent: A Generative Multi-Agent System for Real-Time Industrial Safety Compliance and Predictive Hazard Detection** *Nazia Sultana Plabon*
- 12:15-12:30 **IRAI26-000110 - Feature-Driven Malware Detection Using Random Forests** *Kolade Ademuwagun, Oluwaseyi Ezekiel Olorunshola, Fatimah Adamu-Fika*
- 12:30-12:45 **IRAI26-000257 - Generative Explainable AI for Reliable State Monitoring of Battery Systems** *Muhammad Saeed, Sheroz Shahid Tabassum, Hafiz Muhammad Jawaad, Muhammad Ibrahim, Jiani SHEN, Yijun He*
- 12:45-13:00 **IRAI26-000298 - PHOENIX: Fine-Tuned SLM-Powered Autonomous Satellite Lifetime Extension via Predictive Self-Healing and Multi-Agent AI Recovery** *Sumaiya Islam, Harsha Moraliyage*

📍 501 - Online Only

Papers.

Industrial Applications of AI

- 09:30-09:45 **IRAI26-000200 - Physics Aware Visual Reasoning for Industrial Defect Analysis Using Generative AI** *Govind A, Gadha Lekshmi P, Rohith Syam, Aiswarya Anil Nair*
- 09:45-10:00 **IRAI26-000211 - Low-Resource Domain Adaptation of Small Language Models for Privacy-Preserving IEC 61131-3 Code Generation** *Navneeth Mahadevan, Milan V Raj, MOHAMMED RASHITH K.P*
- 10:00-10:15 **IRAI26-000143 - PilotGPT: A Data Grounded VLM Based AI Copilot for Natural Language Control of Industrial UAVs** *Supun Gamlath, Chathuranga Hettiarachchi, Sulochana Sooriyaarachchi*
- 10:15-10:30 **IRAI26-000229 - MaintainIQ: Multi-Modal Generative AI for Human-Centric Predictive Maintenance in Mining and Manufacturing** *Shamika Dharmasiri, Rakshitha De Silva*
- 10:30-10:45 **IRAI26-000270 - AgentRefine: A Multi-Agent Iterative Refinement Pipeline for Synthetic Defect Generation** *Praapti P, MANJUSRI V, Monicka Jaai, Varshini S, Jagath Shree U*
- 10:45-11:00 **IRAI26-000288 - CircuitForge: A Machine Learning Framework for Constraint-Aware Gate-Level Circuit Synthesis** *Rajesh Bharadwaj, Samarth Pagaku, Sinchana Sunil, Varshini Leopeter*
- 11:00-11:15 **IRAI26-000195 - CounterMaint: Conditional Diffusion Models for Maintenance What-If Trajectory Generation in Industrial Systems** *Harsh Rathva*
- 11:15-11:30 **IRAI26-000154 - AI-Enabled Anomaly Detection for Intelligent Quality Control Using a Hybrid Transformer-VAE Architecture** *Pratyusha Chatterjee, Tanima Bhowmik, sachin kahawala*
- 11:30-11:45 **IRAI26-000285 - TRAFS: A Graph Multi-Agent Reinforcement Learning Framework with Stochastic Scenario Synthesis and Responsible Edge Deployment for Real-Time Urban Traffic Signal Optimization** *Kevin Samuel, Rehaan John, Nithin Joseph F, Adline Sheeba F*
- 11:45-12:00 **IRAI26-000240 - Physics-Informed UV-Vis DSP and Edge One-Class SVM for Real-Time Microbial Contamination Detection in Biomanufacturing** *VASANTH S, Harsha Moraliyage, Swathi R B, Saravanan G*

Saturday 5 September 2026

12:00-12:15 IRAI26-000233 - Perception Pipeline and Edge Deployment of a Power Line Inspection Robot *Saravanan Murugan, David Ekechukwu*

12:15-12:30 IRAI26-000286 - PINN-RL-Generative Framework for Intelligent Pipeline Design Optimization *Kushal Bhargav, Raghav Sandhu, Vinod Sampath, Gaurav Sarma*

12:30-12:45 IRAI26-000373 - Fault Tolerant Industrial Vision Framework Through Semantic-Guided Cross-View Temporal Reconstruction and Restoration *Akshay Gupta, Anshumaan Singh, Rajveer Gupta, Pranshu Jain, Tushar Patidar*

📍 601 - Online Only

Papers.

Industrial Applications of AI - 4

09:30-09:45 IRAI26-000338 - A Planner-Centric Agentic AI Architecture for Execution-Time Production Planning in Energy Equipment Manufacturing *Maheshkumar Patil, Vaibhav Mishra, Prasad Nagella, Manish Kumar Garg*

09:45-10:00 IRAI26-000358 - Benchmarking Edge Computing Platforms: Intelligent Cable Temperature Estimation in Dynamic Line Rating Application *Sahan Siriwardena, Anushka Bandara, Pamudi Abeysinghe, Thanidu Thennakoon, Janaka Ekanayake*

10:00-10:15 IRAI26-000041 - Decoupling Schema Validation in Large Language Models: A Comparative Analysis of Intent Marker Parsers vs. Traditional JSON Function Calling *Amihere Theophilus Junior, Antoniette Nelson Baidoo*

10:15-10:30 IRAI26-000264 - Reactive Control for Discrete Action Vision-Language Navigation: Bridging CoT-VLN Planners and Safe Execution *Lakmina Gamage, Vishwani Geeganage, Sudantha Bopage, Chinthaka Premachandra*

10:30-10:45 IRAI26-000284 - PhysicsDiff-PCB: Physics Guided Conditional Diffusion for Crosstalk Aware PCB Layout Generation *Abishek Jayathilake, Chanaka Madhuranga, Sulitha Bandara, Ramishka Thennakoon, Ushan Kamesh*

10:45-11:00 IRAI26-000183 - The Bilinear Jensen Gap: Fixing a Hidden Failure Mode in Physics-Informed Flow Matching *Devikrupa Bhaskaran, Elizabeth Samuel, Akila Wijethunge*

11:00-11:15 IRAI26-000234 - Co-Partner: A Responsible Agentic AI-Powered Augmented Reality Assistant for Grounded Real-Time Electrical Cable Jointing Guidance *Ammar Helmey Iskandar, Muhd Aqmal Afiq Adnan, Muhammad Khairi Azim, Ezzatul Akmal Kamaru-Zaman*

11:15-11:30 IRAI26-000245 - Generative Scenario Modeling for Robust Dispatch in Renewable-Integrated Smart Grids *Sean Afamefuna, Sadia Rahman, Glory Ojoma Simon, Jessica Osunbor*

11:30-11:45 IRAI26-000231 - Neuro-Diagnostic Digital Twin for Urban EV Infrastructure: A GenAI-Driven Predictive Maintenance Framework *Anthony Ijoot, Roy Muhandi, Ion Miriam Kwagala, Jonathan Serugunda*

11:45-12:00 IRAI26-000147 - Early-Life Prediction of Perovskite Degradation Trajectories: A Benchmark Against Conditional Diffusion Models *Fatemah Lari, Amani Albuloushi, Layaly Salem, Noura AlKhersan, Ghaliah Aldulaimi, Khaled Chahine*

12:00-12:15 IRAI26-000277 - Toward Generative AI-Enhanced Predictive Thermal Management for PV-Based DC Microgrids: A Software-in-the-Loop Study *Utkarsh Sharma, Shrivatsal Sharma, Sneha Narasimhan, Subhashish Bhattacharya*

Saturday 5 September 2026

12:15-12:30 **IRAI26-000149 - Zero-shot Non-Intrusive Load Monitoring** *Anas El Masri, Abdulaziz AlShammari, Ali Ghuluom, Othman Alfaiakawi, Abdulateef Qabazard, Khaled Chahine*

12:30-12:45 **IRAI26-000385 - Cross-Modal Generative Fusion for Interturn Fault Severity Diagnosis in PMSMs** *Fatih Eroglu, Bilal Catkin*

📍 701 - Online Only

Papers.

Responsible AI Applications

09:30-09:45 **IRAI26-000026 - Regularized ResNet-18 Prototypical Networks for Robust Few-Shot Learning for Plant Disease Classification** *Rashadul Islam Sumon, Hayat Ullah, Kiramat Ullah, Hee Cheol Kim, Md Ariful Islam Mozumder*

09:45-10:00 **IRAI26-000228 - MulSenDiff-X: Sensor-Derived Descriptor-Conditioned Diffusion for Unsupervised Multi-Sensor Industrial Anomaly Detection and Root-Cause Explanation** *Benjamin Asamoah Bofo, Jeffery Offei Darko, Shadrach Antoh, Samuel Ankomah Aidoo*

10:00-10:15 **IRAI26-000223 - DocuMind: AI-Native BRD Generation with Traceability, Conflict Detection, and Quality Scoring** *Riddhima Chaturvedi, Nilay Kumar*

10:15-10:30 **IRAI26-000271 - High-Fidelity Synthetic EEG Generation Using 1D U-Net-Based Denoising Diffusion Models for EEG-Centric AI/ML Applications** *Himanshi Virak, Dewashish Lambore, Dhriti Manpurkar, Prasad MVV Kantipudi, Venkateswarlu Gonuguntla*

10:30-10:45 **IRAI26-000273 - Neuro-Sync Adaptive Compression (NSAC): A Multi-Stage Pipeline for LLM Inference at Low Resource Environment** *Saimon Bhuiyan, Md. Kibria Hossen Roni, Mahdiya Rahman Sukanya, Mustakim Bin Mohsin, Mir Sayad*

10:45-11:00 **IRAI26-000263 - MotorGuard AI: Physics-Aware Generative Maintenance with Edge deployment** *Nihaarika PV, Neha Devarasetty, Pujitha Kakarlapudi, Nischal Poudel*

11:00-11:15 **IRAI26-000250 - GenResilience: A Hybrid ML-GenAI Framework for Adaptive Supply Chain Disruption Management** *Lakshya Mehta, Tirth Patel, Aahan Gajera, Raahil Desai, Rutva Patel, Thimal Kempitiya*

11:15-11:30 **IRAI26-000217 - Multimodal Intent Reconstruction for Safety-Critical Communication in Extreme Acoustic Environments** *Roshelle Perera, Naveen Basnayake, Dulavinya Manawadu, Nethmi Pathirana, Manuri Boralugoda, Dahami Fernando*

11:30-11:45 **IRAI26-000022 - Governance-Aware Agent Telemetry for Closed-Loop Enforcement in Multi-Agent AI Systems** *Anshul Pathak, Nishant Jain*

11:45-12:00 **IRAI26-000157 - Uncovering Hidden Neurobiological Patterns in Schizophrenia through Multimodal Neuroimaging and Explainable Clustering** *Nancy Patel, Krish Patel, Jenish Sorathiya, Honey Modha, Himani Trivedi*

12:00-12:15 **IRAI26-000107 - Responsible Synthetic Audience Simulation for Advertisement Pre-Testing: An Evidence-Gated Prototype with Sri Lanka Localisation** *Geemal Muthugala, Thiyoda Pussadeniya, Binura Gamage, Nipun Weerakkody, Thisarani Wickramasinghe*

12:15-12:30 **IRAI26-000348 - AI-Augmented Product Development: A Responsible Methodology for Human-AI Collaboration in Enterprise Product Discovery** *Ganesh Raam Ramadurai, Ganesh Raam Ramadurai*

12:30-12:45 **IRAI26-000085 - Responsible Human-AI Problem Solving: Verification Gaps in a Multinational Wave-1 Pilot Study** *Matthias Huemmer, Theophile Shyiramunda, Franziska Durner, Michelle J.*

Saturday 5 September 2026

Cumings-Koether

10:00-12:30

📍 101

Workshops

Student and YP : Research Connections & Pathways to Industry

This session is for the PhD and Master's students, early-career researchers and young professionals at IRAI, and for Melbourne-based students and researchers from the local IEEE community who want to meet them. The morning has two aims. The first is to get people talking to each other about their research and to find possible collaborations across disciplines and countries. The second is to hear from people who have worked in both academia and industry about how research gets used outside the university, and what that means for your own career.

Program

10:00 Registration and networking. Name tags, coffee, informal introductions.

10:15 Welcome from IEEE YP and SAC. What the two communities do, which universities and organisations are in the room, and how the morning will run.

10:30 Connecting students and researchers through IEEE IAS. Dr Nishad on what the Industry Applications Society offers students and YPs: technical communities, conferences, international networking, volunteering and research opportunities.

11:00 Research lightning introductions. Each participant gets a minute to answer three questions. Who am I? What am I working on? What am I looking for?

11:10 Panel discussion. Dr Nishad, Dr Daniel Marino and Dr Champake Mendis, moderated by Dr Chathurika. How academia, research and industry work together to turn research into something used in practice, how industry collaborations start, what skills matter beyond the PhD, and what career paths exist outside a traditional academic one.

12:00 Rapid fire round. Short questions between participants and panellists.

12:15 Collaboration matchmaking. Participants write down what they can offer and what they are looking for, and find the people whose answers match.

12:30 Closing and next steps. Key takeaways, IEEE opportunities, QR codes, LinkedIn and IEEE connections, group photo.

12:35 Lunch and open networking with international delegates, Melbourne researchers, students, YPs and industry professionals.

IRAI is a research hackathon, and most of the week is spent on the technical problem in front of you. This session is the part of the program that looks past that problem, to the people and industries your work could connect with. Come with something to say about your research and a question for the panel.

Saturday 5 September 2026

	This session will be hosted by the IEEE Victorian Section and IEEE Industrial Applications Society	
12:00-13:00	 Lunch/Coffee Area Lunch	Lunch
13:00-13:30	 101 IRAI Closing Ceremony	Plenary