UQAI
Research Network

2025 UQ AI PhD Showcase

Showcasing AI Research

Hawken Engineering Building (50-C207)
26-27 June 2025

About

UQ's artificial intelligence research network brings researchers together from across disciplines to drive innovation in AI research and its application.

Website: https://ai.uq.edu.au

Email: at uqai@uq.edu.au

Program

Day 1 – Thursday 26 June

Time	Session	Speakers	Page
9:00	**Registration**		
9:30	**Event Opening**		
	Opening address by Pro-Vice-Chancellor (Research)	Prof. Rachael Parker	
	AI Research Network	Prof. Shazia Sadiq	
	Event preview	Dr Alina Bialkowski	
10:00	**Keynote – Building Better Bicycles: The Future of AI**	Dr Peter Bailey (Canva)	5
11:00	**Morning tea**		
11:30	**Presentation Session 1**		
	Anatomical Grounding Pre-training for Medical Phrase Grounding	Wenjun Zhang	8
	Associating cancer-specific single nucleotide variants with methylation changes using models based on convolutional neural networks	Halimat Chisom Atanda	9
	VitalStory: A Human-Centred AI for Capturing Emergency Department Medical Histories	Anton Cush	10
12:00	**Lunch**		
13:00	**Keynote – Beyond the Hype: AI's Role in Shaping a Human-Centred Future**	Lawrence Kusz (ChatStat)	6
14:00	**Poster/demo session and afternoon tea break**		
	AiReview: An Open Platform for Accelerating Systematic Reviews with LLMs	Xinyu Mao	36
	Large Language Models in Biographical Fiction Writing: A Practice-Led Approach	Missy Burrell	37
	Cross-Cohort evaluation of investigating genetic contributions to childhood psychopathology using deep neural networks	Swathi Hassan Gangaraju	38
	Document-level Relation Extraction with Noisy and Large-scale Data	Phan Khai Tran	39
	Robust deep learning models for microwave imaging	Wei-chung Lai	40
	Deep learning to optimize agricultural genomic selection	Chensong Chen	41
	Identification of Important Fragments for Antibiotic Activity using a Deep Learning Model	Abdulmujeeb Onawole	42
	Machine Unlearning for Streaming Forgetting	Shaofei Shen	43
	Converging Systems of Deception: Organisational Risk and Resilience in the Age of LLMs and Deepfakes	Edidiong James	44
	AI Reads the City	Yassin Nooradini	45
	Model-Based Transfer Learning for Traditional ML Models in Clinical Prediction Tasks	Oliver Pienaar	46
	Automated 3D segmentation of glomeruli in human kidney tissue specimens using 16.4 T MRI without contrast agents.	Aurel Amar	47
15:00	**Presentation Session 2**		
	Deep Reinforcement Learning for Robot Control	Humphrey Munn	12
	Image classification with frequency domain visual codes	Wendi Ma	13
	Reassessing Collaborative Writing Theories and Frameworks in the Age of LLMs: What Still Applies and What We Must Leave Behind	Dice Yukita	14
	Enhancing Outbreak Investigations Through a Conversational AI Chatbot	Jovenal Gama-Pinto	15
	Trust vs Disgust: The Believability of Digital Media	Aurora X	16
	Audio-Visual Learning: A Visual Perspective of Sound Localization	Chen Liu	17
	Developing bespoke AI for sport performance and injury prevention: machine learning movement classification in elite women's water polo	Marguerite King	18
16:00-16:15	**Day 1 wrap-up**		

Day 2 – Friday 27 June

Time	Session	Speakers	Page
9:30	**Day 2 Opening**		
9:40	**Presentation session 3**		
	Knowledge Augmented Urban Event Prediction	Fidan Karimova	19
	AI chatbot interventions for symptoms of anxiety and depression, and feelings of loneliness	Laura Neil	20
	Mitigating Language Bias in Academic Research: A Design Science Approach Using Agentic AI	James Boyce	21
	The psychological impact of physically embodied service robots on consumers	David Goyeneche	22
	Towards Scalable and Robust Sign Language Understanding: Datasets, Models, and Applications for Auslan	Xin Shen	23
	Social licence for the secondary use of health data to develop AI	Quita Olsen and Dr Tuan Duong	24
	Forcing Computers to See Fish: Aquaculture and AI	Stella Knief	25
10:40	**Morning tea**		
11:00	**Panel session: AI in the real-world**	Tariq Soliman, Dr Zixin Wang, Dr Xin Xia	7
12:00	**Lunch**		
13:00	**Presentation session 4**		
	Designing Conversational Agents as Trustworthy Mediators for Healthcare Intervention Delivery	Mengyan Hou	11
	The AI Trials: How accommodations and compromises enable momentum in new technology trials	Claire Cunningham	27
	Regularization of Q-learning algorithms for offline reinforcement learning	Jacky Xie	28
	Guiding Patients to Optimal Emergency Facilities: A Computational Design Science Approach for Developing a Large Language Model-Based Tele-Triage System	Hetiao (Slim) Xie	29
	AI as a thinking partner: Exploring international students use of GenAI to support their academic English needs	Franciele Spinelli	30
	Deep learning and genetics for cardiovascular disease prediction: a systematic literature review	Utpala Nanda Chowdhury	31
	AI Generated Scenery for Participatory Urban Design: Application of Crowd-Sourced Feedback and AI-Enhanced Street View Images	Tommy Bao-Nghi Nguyen	32
	Dynamic prediction of ventilation outcomes in critically ill children	Eugene Slaughter	33
	Support Vector Machine in enhancing the diagnosis of Progressive Supranuclear Palsy via Video-oculography	Duy Duan Nguyen	34
	Soft Boundary in Retinal Image Data	Hongwei Sheng	35
	Counting Signatures and Detecting Irregularities: Computer Vision Assisted Election Forensics in Turkey	Ramazan Bora	26
14:20-14:45	**Trivia (with prizes!), closing remarks and presenter awards**		

Keynote Presentations

Building better bicycles: the future of AI

Dr Peter Bailey

Canva

Abstract:
Modern generative AI systems are situated in centuries-old tool building cultures. The advent of modern personal computing created "bicycles for the mind" that accelerated all kinds of digitally enabled experiences. Modern search engines did the same by enabling unparalleled access to digital information. What will AI do in turn? What's holding it back? And why does the user interface experience still feel like we could be talking to ELIZA in 1966?

Biography:
Peter is an expert in the field of information retrieval, and currently leads machine learning for Content & Discovery at Canva. He has worked in a range of organisations, from academia to industrial research labs, from startups and consulting, to more than a decade at Microsoft. He is a co-author of numerous scientific papers and co-inventor of a number of patents. He is the co-creator of several widely used datasets, including WT10g, CERC, UQV100, and CRD3. While at Microsoft, he led the Contextual Relevance team when it introduced personalised search ranking to the Bing web search engine in 2010. Elements of that work were recognised in the ACM SIGIR 2022 Test of Time paper award for "Modeling the impact of short- and long-term behavior on search personalization". In addition to his applied scientist and manager responsibilities, he was a member of Microsoft's Responsible AI AETHER Committee's Fairness and Inclusiveness working group from 2018 through 2021. He is a Senior Member of the Association for Computing Machinery, and contributes regularly to program committees for conferences in the information retrieval community. He is an amateur builder, baker, and marmalade maker.

Beyond the Hype: AI's Role in Shaping a Human-Centred Future

Lawrence Kusz

Founder and Managing Director of Chatstat

Abstract:
In this keynote, I'll reflect on the journey that led me to found Chatstat and work at the intersection of AI, safety, and social impact. We'll explore some of the big-picture questions that today's AI researchers should be considering: What does it mean to be human in an AI-shaped world? How can we ensure AI narrows inequality instead of amplifying it? What are the societal consequences of AGI, personalised learning, AI companions, and synthetic media? And how might AI both challenge and reinforce our sense of identity, truth, and connection?

This won't be a technical deep dive. It is designed to provoke discussion and reflection, especially as you shape the next generation of AI research. Expect a mix of storytelling, provocation, and open questions.

Biography:
Lawrence Kusz is the Founder and Managing Director of Chatstat, an award-winning Australian social enterprise using AI to detect and respond to online harm. He began his career in academia, teaching and researching at the University of Queensland, before founding Chatstat in response to the rise of online risks affecting young people. His work focuses on ethical AI design, digital safety, and the intersection of human behaviour and emerging technology. Lawrence regularly speaks on the societal impacts of AI and has been recognised nationally for his contributions to innovation and social impact.

Panel Discussion

AI in the real-world

Tariq Soliman[1], ***Dr Zixin Wang***[2], ***Dr Xin Xia***[3]



1: Graduate Digital Engineer, Hatch
2: Postdoctoral Research Fellow, ARC Training Centre for Information Resilience (CIRES), UQ
3: Postdoctoral Research Fellow, School of Electrical Engineering & Computer Science, UQ

Abstract:
What does life after graduation look like for those diving into the world of Artificial Intelligence? Join us for an inspiring and candid panel discussion featuring recent UQ graduates who are now working on the frontlines of AI across industry and academia.

This session explores the real-world journey from student to AI professional: where the opportunities are, how to find them, and what skills truly matter beyond the classroom. Our panelists will share how they landed their current roles and reflect on how they apply AI in their work.

We'll tackle topics like the value of academic publications versus a strong online presence, and how much of their studies they actually use on the job. Most importantly, they'll share honest advice they wish they'd known earlier, and tips for today's students and researchers navigating the fast-moving AI landscape.

Whether you're an undergraduate, postgraduate, researcher, or even a staff member looking to give advice to your graduating students, this is your chance to hear real stories, ask questions, and get insight into what it takes to thrive in AI beyond university studies.

List of Presentations

Anatomical Grounding Pre-training for Medical Phrase Grounding

Wenjun Zhang

EECS, EAIT

Presentation abstract: Medical Phrase Grounding (MPG) maps radiological findings described in medical reports to specific regions in medical images. The primary obstacle hindering progress in MPG is the scarcity of annotated data available for training and validation. We propose anatomical grounding as an in-domain pre-training task that aligns anatomical terms with corresponding regions in medical images, leveraging large-scale datasets such as Chest ImaGenome. Our empirical evaluation on MS-CXR demonstrates that anatomical grounding pre-training significantly improves performance in both a zero-shot learning and fine-tuning setting, outperforming state-of-the-art MPG models. Our fine-tuned model achieved state-of-the-art performance on MS-CXR with an mIoU of 61.2, demonstrating the effectiveness of anatomical grounding pre-training for MPG.

Presentation keywords: Vision language model, Medical phrase grounding, Anatomical Grounding

Focus area(s): Data-Centric AI;

Supervisors: Shakes Chandra, Aaron Nicolson

Associating cancer-specific single nucleotide variants with methylation changes using models based on convolutional neural networks

Halimat Chisom Atanda

Mater Research Institute

Presentation abstract: Sequencing technologies have resulted in a relative abundance of genomic data, especially variants, with a limited understanding of their functional impacts. The presence of variants and base modifications on nanopore sequence reads offers an opportunity to link disease-specific variants to methylation changes around affected regions. Here, we present a classification model based on convolutional neural networks to predict somatic variants associated with methylation changes in cancer genomes. We trained the model using phased sequence reads, single-nucleotide germline variants, per-base methylation probability, and variant proximity or presence in differentially methylated regions. Oversampling was applied to account for class imbalance, and we varied model depth, learning rate, and batch size during the training process. The CNN model with the best performance (4 layers, 0.001 learning rate, and batch size 64) had an MCC (Mathews Correlation Coefficient) score of 0.93 (Accuracy = 0.99, Precision = 0.93) in the validation dataset and MCC of 0.94 in the internal test dataset (Accuracy = 0.99, Precision = 0.94). Applying the model produces an output with variant positions, predicted association (positive, negative, or ambiguous), and the direction of methylation change (increase or decrease) with p-values. Our findings suggest that integrating deep learning to cancer genome data may reveal insights into epigenomic biomarkers.

Presentation keywords: AI in cancer, Genomics, Epigenomics

Focus area(s): Data-Centric AI;

Supervisors: Adam Ewing, Sandy Richardson, Yuanhao Yang, Jodi Saunus

VitalStory: A Human-Centred AI for Capturing Emergency Department Medical Histories

Anton Cush

HMBS

Presentation abstract: Emergency departments are high-pressure environments where time constraints and patient stress compromise the quality of clinical communication. VitalStory is a GPT-4.5-powered digital tool developed to capture patient medical histories through conversational AI at the point of arrival in the ED. This tool aims to address the deterioration in memory recall and patient–doctor communication that often occurs with long ED wait times. Drawing on participatory design, Agile development, and Lean UX principles, VitalStory was iteratively developed in collaboration with clinicians and patients at Ipswich Hospital. The forthcoming evaluation phase compares patient-entered histories against structured diagnostic outcomes to assess accuracy, completeness, and clinical usefulness. This project contributes to the responsible integration of generative AI into healthcare, demonstrating how language models can augment clinical workflows and patient agency without replacing the clinician.

Presentation keywords: Human-centred clinical AI

Focus area(s): Human-Centred AI; Scalable and Sustainable AI;

Supervisors: Dr Sisira Edirippulige

Designing Conversational Agents as Trustworthy Mediators for Healthcare Intervention Delivery

Mengyan Hou

EECS

Presentation abstract: What roles could conversational agent play in healthcare intervention deliver when a small number of healthcare professionals are facing 800 patents? How might patients and healthcare professionals say about designing this technology as a mediator rather than an Automator? In this presentation, I will share insights from co-design workshops with both female cancer patients and cancer nurse specialists to reveal the future direction of AI from their lenses.

Presentation keywords: Chatbot, mediator, healthcare, cancer

Focus area(s): Human-Centred AI

Supervisors: Stephen Viller, Chelsea Dobbins

Deep Reinforcement Learning for Robot Control

Humphrey Munn

EECS

Presentation abstract: Full-body robot skills have the potential to transform sectors like farming, manufacturing, and healthcare through advanced automation. My research focuses on reinforcement learning for loco-manipulation robots—those capable of both movement and environmental interaction—by evaluating and developing new methods to improve performance and adaptability on challenging whole-body tasks. I'll discuss my work on full-body robotic throwing, novel techniques I have been working on for enhancing policy adaptability, what I have learnt so far, and reflect on the potential value of this area to UQ's evolving interests in AI and robotics.

Presentation keywords: whole-body control, ppo, reinforcement learning, curriculum learning

Focus area(s): Data-Centric AI;

Supervisors: Marcus Gallagher, Peter Bohm, David Howard, Brendan Tidd

Image classification with frequency domain visual codes

Wendi Ma

EECS

Presentation abstract: Modern computer vision models use mechanisms like convolution and attention to build hierarchical receptive fields (RFs) in the spatial domain. These RFs act as spatial functions defining regions of visual stimuli and expand with model depth. Psychovisual models in the past have instead suggested that human vision uses visual codes, a mechanism that selects coronal sections of the frequency domain, as efficient representations of RFs because convolutions become multiplications in this space, and multiple visual scales can be efficiently accessed in parallel instead of hierarchically. In our work, we introduce PsychoNet, a deep learning model that learns visual codes as an emergent property, providing compelling evidence that they encode RFs. PsychoNet first computes low-level image features in the spatial domain, learns complex-valued representations for them, then uses multi-branch spectral filtering to generate data driven visual codes that perform high-level feature extraction and prediction entirely in the frequency domain. Unlike prior methods, this leverages the spatial and frequency domains in a single, coherent pipeline without going back and forth between them. Experiments on ImageNet classification demonstrate that PsychoNet achieves promising results compared to spatial domain vision models while relying less on model depth.

Presentation keywords: deep learning, architecture, frequency domain

Focus area(s): Data-Centric AI;

Supervisors: Dr. Shekhar Chandra (Principal), AsPr Craig Enstrom (Associate), EmPr Stuart Crozier (Associate)

Reassessing Collaborative Writing Theories and Frameworks in the Age of LLMs: What Still Applies and What We Must Leave Behind

Dice Yukita

EECS

Presentation abstract: In this study, we conduct a critical review of existing theories and frameworks on human-human collaborative writing to assess their relevance to the current human-AI paradigm in professional contexts, and draw seven insights along with design implications for human-AI collaborative writing tools. We found that, as LLMs nudge the writing process more towards an empirical "trial and error" process analogous to prototyping, the non-linear cognitive process of writing will stay the same, but more rigor will be required for revision methodologies. This shift would shed further light on the importance of coherence support, but the large language model (LLM)'s unprecedented semantic capabilities can bring novel approaches to this ongoing challenge. We argue that teamwork-related factors such as group awareness, consensus building and authorship - which have been central in human-human collaborative writing studies - should not apply to the human-AI paradigm due to excessive anthropomorphism. With the LLM's text generation capabilities becoming essentially indistinguishable from human-written ones, we are entering an era where, for the first time in the history of computing, we are engaging in collaborative writing with AI at workplaces on a daily basis. We aim to bring theoretical grounding and practical design guidance to the interaction designs of human-AI collaborative writing, with the goal of enhancing future human-AI writing software.

Presentation keywords: Collaborative writing, large language models, human-AI collaboration

Focus area(s): Human-Centred AI;

Supervisors: Tim Miller, Joel Mackenzie

Enhancing Outbreak Investigations Through a Conversational AI Chatbot

Jovenal Gama-Pinto

SPH

Presentation abstract: Timely and complete data collection is essential for effective outbreak investigations, yet current approaches (e.g. manual contact tracing and food history interviews) are often labour intensive, inconsistent, and slow. The COVID-19 pandemic exposed the limitations of these systems, highlighting the urgent need for scalable digital solutions. Large Language Models (LLMs) offer a promising solution by enabling adaptative, conversational data collection that mimics human dialogue and reduces staff burden. As Australia prepares for the 2032 Brisbane Olympics, enhancing outbreak response capacity is a national priority. The aim of this project is to harness the power of AI to design, develop, and evaluate a chatbot to assist in the collection of case data during disease outbreaks. The objectives include building and fine-tuning a conversational AI chatbot to conduct interviews; evaluating the its usability, and data completeness in simulated or real-world outbreak scenarios; and compare the its data collection against current manual approaches to assess efficiency and quality. This project represents a novel application of AI to a persistent operational barrier in outbreak response. By automating structured interviews, the chatbot could enhance the speed, consistency, and completeness of exposure data, boost surge capacity during health crisis, and reduce reliance on phone-based interviews.

Presentation keywords: Outbreak, data collection, LLM

Focus area(s): Human-Centred AI;AI Strategy and Governance ;Data-Centric AI;Scalable and Sustainable AI;

Supervisors: Amalie Dyda, Elton Lobo, Nicolas Smoll, Martyn Kirk

Trust vs Disgust: The Believability of Digital Media

Aurora X

EECS

Presentation abstract: With the advent of AI-Generated media, the production of synthetic media and so-called deepfakes has become exponentially faster and more accessible. This challenges our understanding of what can really be considered to be "authentic" media online, since anyone with the right tools can produce somewhat convincing footage of speech or events that did not actually occur. How can we deal with the complex arms race of proving the authenticity of media when the believability of synthetic media keeps increasing? In this project we investigate methods we can use to investigate and understand attitudes toward synthetic media technology and its impacts, through a Human Centered Computing perspective.

Presentation keywords: deepfakes, generative AI, synthetic media

Focus area(s): Human-Centred AI

Supervisors: Stephen Viller & Mashhuda Glencross

Audio-Visual Learning: A Visual Perspective of Sound Localization

Chen Liu

EECS

Presentation abstract: Sound waves contain essential information for locating their sources, which is crucial for various applications like navigation, security, and communication. However, accurately identifying the origin of a sound often requires complex analysis, especially in noisy, multiple-sounding source environments or long videos. Due to these challenges, pinpointing exact sounding source localization can be difficult. In this project, I address the challenges of audio-visual segmentation, incorporating multi-modal learning techniques to improve accuracy and efficiency.In the first study year, I first addressed the AVS network degrading into a saliency segmentation network problem due to the data bias problem in existing datasets. To solve the problem, I propose an audio-visual instance-aware segmentation approach. In my second-year study, I aim to establish a reliable semantic mapping between audio and visual signals. To address these issues, I propose a dynamic derivation and elimination framework designed explicitly to mitigate inferior audio-visual alignment caused by intrinsic audio characteristics. Overall, this project develops advanced methods to effectively address the intricate challenge of sounding source localization across various complex scenarios, including handling datasets affected by data bias, managing noise-contaminated audio, achieving robust audio-visual alignment, and mitigating inferior alignment issues arising from inherent audio characteristics.

Presentation keywords: Multi-modal learning, Audio-visual segmentation, Audio Visual Learning

Focus area(s): Human-Centred AI;Data-Centric AI

Supervisors: Xin Yu

Developing bespoke AI for sport performance and injury prevention: machine learning movement classification in elite women's water polo

Marguerite King

School of Health and Rehabilitation Sciences

Presentation abstract: Shoulder injuries have a high incidence and prevalence in women's water polo. Large volumes of overhead movement have been associated with these injuries. There is no measurement method to quantify upper limb movement volumes in women's water polo. A qualitative study was undertaken which indicated that end users (athletes, coaches) required development of tools to measure movement volumes. Project aimed to develop a bespoke machine learning model to classify and quantify overhead movements to assist end users with training load planning. Ten athletes completed four training sessions wearing two inertial measurement units (IMUs) on their dominant wrist and upper back. Sessional video footage and IMU time stamps were synchronised. 19180 movements were coded using custom video analysis for the following movement classifications: swimming, high intensity and low intensity throwing and blocking with and without ball contact. Raw IMU features were then computed and synchronised with coded video movement classification which served as ground truth. IMU features were computed then classification random forest model developed using a 75:25 training-test partition. Class imbalances were accounted for using the Synthetic-Minority-Oversampling-Technique (SMOTE). Model was optimised using a 15 k-fold cross validation. Model balanced accuracy was 99.3%. This model was then applied to predict movement classifications across a water polo season for coaches and athletes.

Presentation keywords: machine learning; water polo; inertial measurement unit; injury prevention

Focus area(s): Human-Centred AI

Supervisors: Professor Bill Vicenzino, Dr Raimundo Sanchez, Dr Nathalia Costa, Dr Kate Watson, Associate Professor Michelle Smith

Knowledge Augmented Urban Event Prediction

Fidan Karimova

EECS

Presentation abstract: Urban event prediction remains a critical yet underdeveloped area in spatiotemporal machine learning, particularly when it comes to generalizability, contextual relevance, and adaptability. This PhD research addresses three major gaps: (1) the lack of cross-city generalization due to inconsistent event categorizations and heterogeneous dynamics; (2) the absence of contextual awareness in existing models, which ignore societal signals such as policy shifts or public events; and (3) the inability of current methods to adapt to rapid, non-stationary changes in urban patterns. To tackle these challenges, we first developed HYSTL, a hypernetwork-based framework that learns from multiple cities with mismatched event types by leveraging a structured knowledge graph to guide prediction. Building on this foundation, future work will incorporate real-time societal signals (e.g., news, policies) to improve interpretability and predictive responsiveness, and develop lightweight, adaptive models that maintain performance under dynamic conditions without costly retraining. Collectively, this research advances the field toward more robust, transferable, and context-aware event prediction systems.

Presentation keywords: knowledge graph, urban event prediction

Focus area(s): Data-Centric AI

Supervisors: Dr. Rocky Chen, Prof. Shazia Sadiq

AI chatbot interventions for symptoms of anxiety and depression, and feelings of loneliness

Laura Neil

School of Pharmacy & Pharmaceutical Sciences

Presentation abstract: Title: The effectiveness of AI chatbots in treating symptoms of anxiety and depression, and feelings of loneliness

Objective: The aim of this review was to assess the effectiveness of AI chatbots in treating symptoms of anxiety and depression, and feelings of loneliness.

Methods: This systematic review was conducted using PubMed, CINAHL, and Embase. Key search terms included artificial intelligence, chatbots, mental health, anxiety, depression, loneliness, conversational agent and large language model. Studies were included if they were published in peer-reviewed journals after 2022, involved AI chatbots that are synchronous, two-way and text based, and used a psychometric scale to determine mental health outcomes for effectiveness measurement.

Results: There were 65 articles included in this review. Some of the most commonly used scales were the GAD-7 for anxiety, PHQ-8 for depression, and the UCLA for loneliness. Most studies used psychometric scales at baseline and post-intervention, although some compared psychometric scale results of participants who used AI chatbots to those who underwent usual care or an alternative comparator such as an online psychoeducation course. Results of AI chatbot use appear to be equivalent to usual care or improved user outcomes.

Conclusion: This review demonstrates that AI chatbot interventions have potential to reduce symptoms of anxiety or depression, and feelings of loneliness.

Presentation keywords: AI Chatbots, Mental Health, Anxiety, Depression, Loneliness

Focus area(s): Human-Centred AI

Supervisors: Dr. Centaine Snoswell, Prof. Liam Caffery, Dr. Helen Haydon, Dr. Aaron Snoswell

Mitigating Language Bias in Academic Research: A Design Science Approach Using Agentic AI

James Boyce

BEL

Presentation abstract: Language bias in academic publishing creates systemic disadvantages for non-native English researchers. Beyond translation issues, it reinforces deeper epistemic, stylistic, and behavioural norms aligned with English-centric academic traditions and gatekeeping, such as publication, visibility, and global collaboration. Researchers must not only translate their work but adapt to rhetorical and disciplinary expectations shaped by English-language norms. This thesis explores language bias through a socio-technical lens, using the Echelon Design Science Research (eDSR) methodology to design and evaluate three interlinked AI-driven artefacts. Phase 1 introduces a Translator tool for high-fidelity, domain-sensitive academic translation via a multi-agent LLM framework. Phase 2 proposes an agentic Lay Summarizer to enhance accessibility and cross-disciplinary engagement through audience-aware summaries. Phase 3 plans to integrates both into an Inclusive Support System for inclusive writing, reviewing, and reading. Each artifact is iteratively developed and evaluated via testing, user feedback, and real-world use. Early results from Phase 1 show strong preservation of meaning, tone, and terminology. Upcoming evaluations will examine broader impacts on comprehension, accessibility, and inclusive academic practice. This research contributes novel multilingual support tools and advances eDSR to tackle one of academia's most persistent equity challenges: language bias.

Presentation keywords: Translation, Multilingual research, Large Language Models (LLMs), Multi-agent systems, Language bias

Focus area(s): Human-Centred AI

Supervisors: Morteza Namvar, Saeed Akhlaghpour, Andrew Burton-Jones

The psychological impact of physically embodied service robots on consumers

David Goyeneche

marketing, Business School

Presentation abstract: This presentation draws on my PhD research in Marketing, which examines how physically embodied service robots influence consumer self-perception in service encounters. While most research emphasizes robots' technical functions or acceptance (Leung et al., 2018; Van Osselaer et al., 2020), this work explores how consumers psychologically experience and interpret robots during service interactions.

The thesis includes three interconnected manuscripts. The first is a systematic literature review that organizes service robot theories into eight themes and introduces a triadic framework linking consumers, employees, and firms (De Keyser and Kunz, 2022; Wirtz et al., 2018). The second and third manuscripts are supported by a total of ten experimental studies. The second shows that interacting with service robots can enhance consumers' feelings of competence, particularly under uncertainty (Ryan and Deci, 2000; De Bellis and Johar, 2020). The third finds that robots, compared to humans, provide less social validation in preferential treatment contexts, reducing perceived status and gratitude (Drèze and Nunes, 2009; Palmatier et al., 2009).

Together, the manuscripts identify self-perception as a key mechanism in consumer–robot interactions and offer practical insights for organizations aiming to align technological innovation with psychological and social needs (Chiang and Trimi, 2020; Shanks et al., 2024).

Presentation keywords: Robot, AI, Self-perception

Focus area(s): Human-Centred AI

Supervisors: Nicolas Pontes and Peter Popkowski

Towards Scalable and Robust Sign Language Understanding: Datasets, Models, and Applications for Auslan

Xin Shen

EECS

Presentation abstract: This presentation introduces a comprehensive pipeline for advancing Australian Sign Language (Auslan) technology, covering data, modeling, and deployment. We first present MM-WLAuslan, a large-scale, multi-view, multi-modal dataset for isolated Auslan recognition. Then, we propose a cross-view learning framework that synthesizes novel viewpoints and disentangles semantic and view-specific features, improving robustness under camera variation. Finally, we introduce AuslanWeb, a web-based, bidirectional Auslan-English translation system powered by large language models. It supports both isolated and successive signing and has been validated through user studies. Together, these efforts advance sign language research and practical accessibility in real-world scenarios.

Presentation keywords: Auslan, Sign Language Recognition, Communication System

Focus area(s): Human-Centred AI; Data-Centric AI

Supervisors: Dr Xin Yu, Prof Helen Huang

Social licence for the secondary use of health data to develop AI

Dr Tuan Duong and Quita Olsen

Queensland Digital Health Center, Centre for Health Services Research

Presentation abstract: Introduction: Artificial intelligence (AI) in healthcare holds the potential to improve health care. However, using personal health data to train AI models requires a social licence. To address this gap, we conducted a systematic exploration of global attitudes toward sharing health data for AI development and will also examine consumer perspectives on data use and offer recommendations to inform best practice guidelines aimed at strengthening the social license for AI integration in healthcare. A meta-analysis was conducted to examine worldwide willingness to share health data for secondary purposes. In addition, qualitative research will be undertaken encouraging participants to reflect on the use of personal health data for training AI models. There were 65 eligible studies for willingness to share and 6 studies specific for sharing health data for AI. The global willingness to share health data estimate was 77% and willingness to share for AI purposes was 76% in the random effects models. The systematic review showed participants remained concerned about privacy, consent, and transparency. Future work should consider public education program and assessing diverse populations. The findings from the qualitative research will further provide key consumer concerns and provide recommendations to guide best practices for establishing a robust social license for AI integration in healthcare.

Presentation keywords: Social licence, consent, transparency

Focus area(s): AI Strategy and Governance

Supervisors: Prof Clair Sullivan, Prof Jason Pole, Dr Lee Woods, Dr Amalie Dyda

Forcing Computers to See Fish: Aquaculture and AI

Stella Knief

SENV

Presentation abstract: Computer vision (CV) is a powerful yet currently underutilized tool in aquaculture. The slow adoption by farmers may be due to the perceived complexity of developing a CV model, the anticipated high infrastructure demands, as well as lack of awareness of the potential benefits of working with CV technology. I herein address these concerns and briefly present two of my own works in this space. First, a case study in developing an accessible and effective CV model capable of detecting stress behaviours in barramundi for an Australian barramundi farm. The model was trained on farm footage and developed with the intent to act as a "early warning" system for the farmers when temperature stress or other environmental stress would impact the welfare and overall health of their fish. Second, I present my ongoing work towards a model capable of detecting disease in rainbow trout, and where this work fits in the broader context of disease-focused CV work in aquaculture. Lastly, I direct our gaze to the future and share an outlook on the future of CV, and artificial intelligence at large in aquaculture.

Presentation keywords: computer vision, aquaculture, fish health, farm management

Focus area(s): Human-Centred AI

Supervisors: Prof. Andrew Barnes, Dr. Rebecca Cramp

Counting Signatures and Detecting Irregularities: Computer Vision Assisted Election Forensics in Turkey

Ramazan Bora

School of Economics

Presentation abstract: Election integrity in semi-functional democracies is an important factor for the credibility of democratic institutions. In the context of Turkey, I examine the role of semi-partisan Ballot Box Committees (BBCs) in administering fair elections, specifically, the relationship between invalid votes and BBC size. Because structured digital records on BBC attendance do not exist, I develop a computer-vision pipeline that processes millions of handwritten, signed ballot-box reports from the twin 2015 parliamentary elections and the 2023 presidential elections. A custom-trained YOLOv11 model detects member signatures in these reports. Multiple signature sources and strict confidence thresholds are used to cross-validate counts and yield a high-quality panel of BBC sizes at the ballot-box level. Using a two-way fixed-effects identification strategy that exploits within-ballot-box (and, when necessary, within-neighbourhood) variation in BBC size across election cycles, I estimate a precise null: additional BBC members do not change invalid-vote shares. These results align with recent Latin-American evidence, underscoring the limited role of election observers in influencing ballot rejections in modern elections.

Presentation keywords: computer vision, yolo, signature detection

Focus area(s): Human-Centred AI

Supervisors: Prof Alicia Rambaldi, Dr Haishan Yuan

The AI Trials: How accommodations and compromises enable momentum in new technology trials

Claire Cunningham

Business School, CIBIT

Presentation abstract: Artificial Intelligence (AI) is reshaping organizational practices but poses challenges when individuals, teams, and organizations hold differing perceptions of the technology. These challenges are amplified in multi-organizational collaborations. While technological framing theory (Orlikowski & Gash, 1994) has explored the effects of frame incongruence on technology deployment, less is known about how congruent frames are maintained during implementation. This gap assumes that congruence ensures smooth, successful trials, overlooking the work required to maintain it.

My study investigates how technological frame congruence was maintained during a 29-month ethnographic trial of an AI mammography tool in Vietnam, involving a Trial Facilitator, Technology Provider, and Technology User. Using interviews, observations, and document analysis, I examined technological frames across three domains: Nature of Technology, Technology Strategy, and Technology-in-Use. Thirteen episodes revealed that congruence was preserved through two mechanisms: compromises and accommodations.

I develop a process model showing how these micro-level framing practices help navigate challenges such as communication issues, resource allocation, and sustaining trial momentum. The findings offer practical insights into managing collaborative technology trials and highlight the active work required to maintain frame alignment for successful implementation.

Presentation keywords: Technological frames, New technology trials

Focus area(s): AI Strategy and Governance

Supervisors: Associate Professor Anna Jenkins, Associate Professor Sam MacAulay

Regularization of Q-learning algorithms for offline reinforcement learning

Jacky Xie

SMP

Presentation abstract: Offline reinforcement learning is the problem of learning optimal decision-making strategies from a fixed dataset, without further exploration. Compared to traditional online reinforcement learning, classical algorithms such as Q-learning, when applied to offline datasets, suffer from out-of-distribution queries that cause extrapolation errors during training. In this seminar I will present research that focuses on regularizing the Q-learning algorithm in a pessimistic way by adding a density-based penalty to the learning objective.

Presentation keywords: Reinforcement learning, machine learning

Focus area(s): Data-Centric AI; Scalable and Sustainable AI

Supervisors: Nan Ye, Fred Roosta

Guiding Patients to Optimal Emergency Facilities: A Computational Design Science Approach for Developing a Large Language Model-Based Tele-Triage System

Hetiao (Slim) Xie

Business School

Presentation abstract: The recent establishment of Minor Illness and Injury Clinics (MIIC) in Australia provides a pathway for low-acuity patients to receive appropriate urgent care and for improving healthcare system efficiency. However, the lack of a navigation system prevents this program from being utilized effectively, as patients struggle to select the emergency facility best suited to their clinical needs. Recent advancements in artificial intelligence are revolutionizing healthcare and highlighting opportunities to develop advanced AI models to address this challenge. Therefore, we follow the computational design science approach to propose a large language model (LLM)-based tele-triage system. This computational artifact takes into account the salient characteristics of the problem drawn from both problem-related and general machine learning (ML) literature, as well as from the interview analysis with clinical experts. Based on the formulated problem, we will design three sequential and interconnected novel LLM-based modules that use constructed synthetic patient self-reported data to generate predictions on whether patients should attend the emergency department or MIIC. Comprehensive computational and field evaluations will then demonstrate the novelty and validity of the design. We expect our proposed artifact to make significant methodological contributions to the knowledge base in both information systems and digital health fields, and to improve patient care.

Presentation keywords: large language models, healthcare predictive analytics, digital health

Focus area(s): Human-Centred AI

Supervisors: Dr Morteza Namvar, A/P Saeed Akhlaghpour, Prof Marten Risius, Prof Andrew Burton-Jones, A/P Andrew Staib

AI as a thinking partner: Exploring international students' use of GenAI to support their academic English needs

Franciele Spinelli

HASS, School of Languages and Cultures

Presentation abstract: Despite meeting university English entry thresholds, many international students who speak English as a second language (L2) face academic language challenges, often adopting technology to support their reading, writing, speaking and listening needs. GenAI tools such as ChatGPT have expanded international students' academic toolkit, acting as a thinking partner to help them meet the demands of academic English. Yet, while GenAI use requires self-regulated learning (SRL) skills that enable learners to manage and assess their learning, few L2 studies have examined this cohort's use of GenAI through this theoretical lens. Drawing on key SRL constructs, this mixed-methods study explores how and why international students at an Australian university utilise GenAI for academic language support. Preliminary results from 150 questionnaires and 28 semi-structured interviews show that international students use GenAI tools to improve their grammar and clarity in writing and speaking, simplify complex readings and confirm listening comprehension. Some of their motivations for using this technology are managing cognitive load and increasing confidence in their English skills. While many critically engage with GenAI outputs, some are concerned that unbalanced use of the technology could hinder their English development. Findings highlight the need for scaffolded support and curriculum-embedded GenAI practices to help international students self-regulate and succeed academically.

Presentation keywords: GenAI; International students; Academic English; Self-regulated learning; Human-centred AI.

Focus area(s): Human-Centred AI

Supervisors: Dr Peter Crosthwaite and Dr Simone Smala

Deep learning and genetics for cardiovascular disease prediction: a systematic literature review

Utpala Nanda Chowdhury

IMB

Presentation abstract: The genetic basis of cardiovascular disease (CVD) has led to use deep learning (DL) for CVD prediction from genetic data. However, their effectiveness and challenges remain unclear. This systematic review examines recent advances (2019-2024) in applying DL for CVD prediction using genetic data. A PubMed and Scopus search selected 56 studies from 1,109 abstracts using PRISMA guidelines. A wide range of DL models was applied to genetic data alone or combined with cardiac imaging (MRI), electrocardiogram (ECG) signals, and clinical and demographic variables. Atrial fibrillation was the most frequently studied CVD, with most studies extracting novel features linked to CVD from MRI and ECG. Other applications included genotype prediction, risk stratification, improving polygenic risk scores, and guiding variant analysis. Transfer learning with pretrained U-Net and CNNs (ResNet8/34/50, VGG-16) was commonly used for image segmentation, with a ResNet34 model reaching the highest dice-score of 0.97. CNNs attained accuracies up to 0.98 in some tasks, while a ResNet-based model yielded an AUC of 0.99 for variant pathogenicity prediction. Reproducibility was a concern, as over half of the studies did not share codes/model weights. Explainability was addressed in only one-third cases, and the dominance of European population limited the generalizability. We recommend a focus on interpretable models, transparent reporting, model sharing and external validation in future.

Presentation keywords: Deep learning, cardiovascular disease, genetics, review, prediction

Focus area(s): Data-Centric AI

Supervisors: Dr Sonia Shah, Dr Baptiste Couvy-Duchesne

AI Generated Scenery for Participatory Urban Design: Application of Crowd-Sourced Feedback and AI-Enhanced Street View Images

Tommy Bao-Nghi Nguyen

School of Architecture, Design and Planning

Presentation abstract: The formulation and design of urban spaces are essential to our health as a society and are well documented in literature and practice. However, existing approaches fall along a spectrum between qualitative design methods (surveys, design workshops) and quantitative methods using big data analytics. While quantitative design methods provide designers and policymakers with large volumes of precise data to inform decision-makers, they often struggle with the underlying 'why' and 'how' behind user behaviour. Conversely, qualitative methods provide rich, nuanced understanding but are time-consuming and resource-intensive, limiting their scalability.

Generative AI is an emerging tool in urban design, and this research explores the potential to generate alternative streetscape scenarios that help reveal public preferences and better understand what the public values in urban environments. With the Olympics approaching, this thesis explores Brisbane as a case study. This research applies its method to evaluate how focused street designs can be more efficiently developed using generative AI. The outcome is a scalable, perception-informed framework contributing to methodological advancement and practical urban design.

Presentation keywords: Generative AI, Urban Design, Participatory Design, Street View Images

Focus area(s): AI Strategy and Governance

Supervisors: Dr Dan Luo, Dr Silvia Micheli, Prof. Ali Cheshmehzangi

Dynamic prediction of ventilation outcomes in critically ill children

Eugene Slaughter

Faculty of Health, Medicine, and Behavioural Sciences

Presentation abstract:

Background: Invasive mechanical ventilation (IMV) is a vital intervention in paediatric intensive care units (PICUs), however when unplanned can lead to complications associated with prolonged PICU stay and mortality. IMV also poses risks if withdrawn prematurely, potentially requiring re-initiation, known as failed extubation (FE). Early warning tools for IMV requirement and FE may enable planning or prevention, improving clinical outcomes.

Aim: To develop predictive models for IMV requirement and FE in PICU patients.

Methods: EHR data from 28,066 patients admitted to two Australian PICUs from 2015 to 2024, and traditional machine learning techniques, including logistic regression, random forest, and gradient boosted trees (GBT) were explored to develop real-time and on-demand prediction models for IMV requirement and FE, respectively.

Results: A GBT model achieved optimal performance for the prediction of IMV requirement, with good discriminative performance (AUROC: 0.84 [0.80, 0.87]). Salient features included admission source, fraction of inspired oxygen, and previous IMV length. To-date, a performant model for FE has not been identified, with best performing models achieving relatively poor discriminative performance (AUROC: 0.63 [0.55, 0.68]).

Conclusion: Our preliminary findings indicate promise for the development of explainable and performant ML-based early warning tools for ventilation outcomes in PICU patients. Further refinements and evaluation however are required.

Presentation keywords: paediatric ICU ML ventilation prediction

Focus area(s): Human-Centred AI

Supervisors: Prof. Kristen Gibbons, Dr. Sainath Raman, Dr. Patricia Gilholm

Support Vector Machine in enhancing the diagnosis of Progressive Supranuclear Palsy via Video-oculography

Duy Duan Nguyen

Queensland Brain Institute, Mater Research Institute, HMBS, UQ

Presentation abstract: Progressive Supranuclear Palsy (PSP) is a neurodegenerative disorder often misdiagnosed as Lewy body disorders (LBD) or Alzheimer's disease (AD) due to overlapping symptoms. Current clinical assessments are subjective and limited in early-stage detection. This project applies Support Vector Machine (SVM) models to video-oculography (VOG) data to improve diagnostic accuracy using objective eye movement biomarkers.

We collected VOG data from 92 patients (21 PSP, 50 LBD, 21 AD) and 82 controls using a series of oculomotor tasks (saccades, smooth pursuit, and square wave jerk detection). Key features included vertical saccade amplitude, misaligned saccadic angle, and smooth pursuit velocity accuracy. Classification performance was assessed using AUC and SVM models.

Vertical saccade amplitude showed the best single-feature AUC (0.97 for PSP vs. controls). Novel features—misaligned angle and pursuit accuracy—also provided strong discrimination. Combining all features in an SVM model yielded high classification accuracy (AUC = 0.94 vs. controls, 0.90 vs. LBD, 0.86 vs. AD).

This study demonstrates the potential of AI-driven approaches to enhance clinical diagnosis. The SVM-based model offers a reproducible, non-invasive and stable method to differentiate PSP from other neurodegenerative diseases with high accuracy.

Presentation keywords: Video-oculograph, Progressive Supranuclear Palsy, Support Vector Machine

Focus area(s): Data-Centric AI

Supervisors: Prof. Peter Nestor

Soft Boundary in Retinal Image Data

Hongwei Sheng

EECS

Presentation abstract: Retinal imaging plays a central role in ophthalmological deep learning research due to its ability to reveal detailed information about eye and systemic health. It supports automated detection of diseases such as diabetic retinopathy and glaucoma. A key topic in this field is the soft boundary problem, where disease categories lack clear distinctions because of biological variability and overlapping features. This ambiguity makes it difficult to train accurate and reliable models for clinical use. In this work, we briefly analyze the presence and impact of soft boundaries in retinal imaging, highlighting their implications for model development and clinical reliability.

Presentation keywords: Retina Imgaing, Computer Vision, Deep Learning

Focus area(s): Data-Centric AI

Supervisors: Dr. Xin Yu

List of Posters

AiReview: An Open Platform for Accelerating Systematic Reviews with LLMs

Xinyu Mao

EECS

Demo abstract: Systematic reviews are fundamental to evidence-based medicine. Creating one is time-consuming and labour-intensive, mainly due to the need to screen, or assess, many studies for inclusion in the review. Several tools have been developed to streamline this process, mostly relying on traditional machine learning methods. Large language models (LLMs) have shown potential in further accelerating the screening process. However, no tool currently allows end users to directly leverage LLMs for screening or facilitates systematic and transparent usage of LLM-assisted screening methods. We introduce (i) an extensible framework for applying LLMs to systematic review tasks, particularly title and abstract screening, and (ii) a web-based interface for LLM-assisted screening. Together, these elements form AiReview-a novel platform for LLM-assisted systematic review creation. AiReview is the first of its kind to bridge the gap between cutting-edge LLM-assisted screening methods and those that create medical systematic reviews.

Demo keywords: Large Language Model, Systematic Review, Title and Abstract Screening

Focus area(s): Human-Centred AI; Data-Centric AI;

Supervisors: Guido Zuccon, Bevan Koopman, Harrisen Scells

Large Language Models in Biographical Fiction Writing: A Practice-Led Approach

Missy Burrell

School of Communications and Arts

Poster abstract: This research explores how large language models (LLMs) can assist writers during the research and drafting phases of biographical fiction, without replacing the creative process central to authorship. While LLMs are capable of generating human-like text, this practice-led thesis investigates their use as supportive tools that preserve the writer's artistic control and narrative integrity. Through experimentation with a locally hosted LLM, I examine applications such as summarising large data volumes, referencing with RAG-style prompting, offering creative prompts, and revising drafts, while ensuring the author remains the primary creative force. A reflective journal will document challenges, insights, and decision-making throughout the writing of a biographical fiction novel. This project addresses a gap in current research by providing a critical framework for ethical, effective, and author-led use of GenAI in literary practice, offering guidance for writers who seek to responsibly incorporate AI without compromising originality or voice.

Poster keywords: 'Large Language Models' 'Biographical Fiction'

Focus area(s): Human-Centred AI;

Supervisors: Leah Henrickson, Kim Wilkins

Cross-Cohort evaluation of investigating genetic contributions to childhood psychopathology using deep neural networks

Swathi Hassan Gangaraju

Child Health Research Centre

Poster abstract: At the individual level, genetic susceptibility to a psychiatric disorder can be conceptualised as comprising two components: a central core representing disorder-specific genetic risk, and a surrounding layer that reflects shared genetic risk with other comorbid disorders and traits. Whether incorporating both shared and unique genetic risks to approximate an individual's overall genetic predisposition can meaningfully improve predictive models for specific disorders remains unexplored. Hence, we aim to investigate the predictive utility of the Polygenic risk score (PRS) by integrating the PRS of a target disorder with those of comorbid traits. Our baseline models will include logistic regression, and the main model will be a deep neural network, which has the potential to handle high-dimensional polygenic inputs. This study utilises three well-established longitudinal cohorts: the UK Millennium Cohort Study (UK-MCS), the Adolescent Brain Cognitive Development (ABCD) study, and the Mater-University of Queensland Study of Pregnancy (MUSP). For the cross-cohort analysis, the UK-MCS cohort will be used to train the models, while the ABCD and MUSP cohorts will serve as independent test sets to evaluate model generalizability, which examines two outcomes: Depression and ADHD. This integrative approach is designed to enhance our understanding of the genetic architecture underlying adolescent psychiatric disorders and to improve the prediction of psychiatric outcomes.

Poster keywords: polygenic risk scores, shared genetic liabilities, comorbidity, deep neural network

Focus area(s): Data-Centric AI;

Supervisors: Dr Enda Byrne, Prof Christel Middeldorp

Document-level Relation Extraction with Noisy and Large-scale Data

Phan Khai Tran

School of EECS

Poster abstract: The explosion of unstructured text from the internet and mobile platforms presents a key challenge: extracting high-quality structured information from noisy, diverse sources. Document-level Relation Extraction (DocRE), which identifies relationships across multiple sentences, supports downstream tasks like knowledge graph construction and Large Language Models based Question Answering systems. Yet, DocRE suffers from data-centric bottlenecks: (1) noisy and irrelevant content, (2) imbalanced relation distributions, and (3) limited labeled data due to annotation costs.

We approach DocRE from a data-centric AI perspective, focusing on improving data quality and utility. First, we propose an evidence sentence selection framework using a dynamic bipartite graph and attention mechanism to filter irrelevant content and enhance relational context. Second, we address long-tail relation imbalance via embedding-space data augmentation, generating relation-specific entity pair representations. Third, we design a semi-supervised learning framework with co-training teacher-student models to create reliable pseudo-labels from weakly labeled and unlabeled data — without extra training overhead. These contributions demonstrate how better data handling can enhance DocRE's robustness and real-world applicability.

Poster keywords: Natural Language Processing, Relation Extraction, Deep Learning

Focus area(s): Data-Centric AI

Supervisors: Professor Xue Li, Associate Professor Wen Hua

Robust deep learning models for microwave imaging

Wei-chung Lai

EECS

Poster abstract: Microwave imaging (MWI) is a non-ionising, cost-effective, and portable alternative to conventional medical imaging. By exploiting dielectric contrasts within biological tissue, MWI reconstructs internal structures from scattered electromagnetic signals. While MWI shows potential in applications such as stroke detection, traditional reconstruction techniques often fail in complex, heterogeneous environments and lack the speed required for emergency use. To address these limitations, deep learning has gained increasing interest in MWI, driven by its success in other medical imaging domains. However, the adoption of deep learning in MWI faces three key challenges, including limited availability of high-quality data, interference from dominant clutter signals, and the inherently low interpretability of deep learning models. This presentation explores deep learning strategies designed to address the key challenges in applying deep learning to MWI. Specifically, it demonstrates how generative models can augment training data, how self-supervised learning can mitigate clutter without relying on paired datasets, and how explainability techniques can improve the transparency of model predictions by incorporating physical intuition. In summary, these approaches aim to advance MWI toward more practical and reliable clinical solutions.

Poster keywords: Deep learning, microwave imaging, stroke detection

Focus area(s): Data-Centric AI

Supervisors: Alina Bialkowski, Lei Guo

Deep learning to optimize agricultural genomic selection

Chensong Chen

QAAFI

Poster abstract: In crop and livestock breeding, combining numerous desirable alleles for important agronomic traits, such as disease resistance, drought and heat tolerance, end-use quality and high yield potential in the shortest possible time remains a major challenge. Because of the number of alleles that need to be considered for these traits, the number of possible mating combinations increases exponentially. Deep learning technologies have shown their great potentials in capturing hierarchical information from complex architecture, were considered to address current challenges in genomic prediction to capture nonlinear complex genetic architectures and increase the prediction accuracy. We hereby introduce a hybrid framework that incorporated attention mechanism and external Bayesian-based marker selection that successfully enhance the prediction outcomes in sugarcane disease prediction, which was identified as a very complex strand of traits.

Poster keywords: Genomic selection, Genomic prediction, Deep learning, Quantitative genetics

Focus area(s): Data-Centric AI; AI Strategy and Governance

Supervisors: Ben Hayes

Identification of Important Fragments for Antibiotic Activity using a Deep Learning Model

Abdulmujeeb Onawole

Institute for Moleclar Bioscience

Poster abstract: Antimicrobial resistance threatens global health, requiring new approaches to drug discovery. While machine learning can predict antibiotic activity with high accuracy, "black box" models provide limited insights for drug design. We developed an explainable Relational Graph Convolutional Network (RGCN) to identify molecular features driving activity against Staphylococcus aureus. Our model achieved excellent performance (AUC: 0.909, Precision: 0.938) on 54,277 S. aureus tested compounds. Fragment attribution analysis identified 105,077 positively contributing molecular fragments (attribution score >0.1), including 64,632 scaffolds and 40,445 substituents, with average attribution scores of 0.59 and 0.27 respectively. Attribution scores ranged from 0.10 to 0.76, with a median of 0.47. This interpretable framework enables virtual screening by prioritizing compounds containing high attribution fragments and supports rational design by combining favourable molecular scaffolds. The approach demonstrates how explainable AI can transform computational predictions into actionable chemical insights for antimicrobial discovery.

Poster keywords: Deep Learning; Explainabl AI; Antimicrobial resistance; Drug Discovery

Focus area(s): Data-Centric AI

Supervisors: Dr. Johannes Zuegg and Prof. Mark Blaskovich

Machine Unlearning for Streaming Forgetting

Shaofei Shen

EECS

Poster abstract: Machine unlearning aims to remove knowledge of the specific training data in a well-trained model. Currently, machine unlearning methods typically handle all forgetting data in a single batch, removing the corresponding knowledge all at once upon request. However, in practical scenarios, requests for data removal often arise in a streaming manner rather than in a single batch, leading to reduced efficiency and effectiveness in existing methods. Such challenges of streaming forgetting have not been the focus of much research. In this paper, to address the challenges of performance maintenance, efficiency, and data access brought about by streaming unlearning requests, we introduce a streaming unlearning paradigm, formalizing the unlearning as a distribution shift problem. We then estimate the altered distribution and propose a novel streaming unlearning algorithm to achieve efficient streaming forgetting without requiring access to the original training data. Theoretical analyses confirm an $O(\sqrt{T}+V_T)$ error bound on the streaming unlearning regret, where V_T represents the cumulative total variation in the optimal solution over T learning rounds. This theoretical guarantee is achieved under mild conditions without the strong restriction of convex loss function. Experiments across various models and datasets validate the performance of our proposed method.

Poster keywords: Machine Unlearning

Focus area(s): Data-Centric AI

Supervisors: Miao Xu, Weitong Chen, Alina Bialkowski

Converging Systems of Deception: Organisational Risk and Resilience in the Age of LLMs and Deepfakes

Edidiong James

UQ Business School

Poster abstract: This presentation explores how the convergence of large language models (LLMs) and deepfake technologies is enabling a new class of AI-driven deception that challenges existing models of organisational security and digital trust. While LLMs and synthetic media are often examined independently, their integration creates compounded risks that remain under-addressed in current research and practice. This work forms part of my broader PhD project focused on identifying and addressing organisational vulnerabilities arising from generative AI convergence. Through a multi-phase investigation, the research examines how adversaries combine synthetic text, voice, and video to craft persuasive attacks, drawing on findings from a systematic literature review and expert workshops with cybersecurity professionals across Australia and the UK. The showcase will highlight key gaps in cross-modal detection, institutional response planning, and regulatory governance. The research also engages criminological and AI policy frameworks to examine institutional readiness and regulatory blind spots. Early-stage concepts for a socio-technical prototype will be presented, synthesising insights across the studies to support scenario-based threat assessment and strategic planning. By bridging empirical, technical, and policy domains, this work offers an interdisciplinary contribution to AI safety, strategy, and organisational resilience in an era of converged synthetic threats.

Poster keywords: LLMs, Deepfakes, Generative AI, Organisational Security, AI Governance

Focus area(s): AI Strategy and Governance ;Human-Centred AI

Supervisors: Dr Ivano Bongiovanni (UQ Business School), A/Prof Mashhuda Glencross (EECS, UQ) , Dr Lewys Brace (University of Exeter)

AI Reads the City

Yassin Nooradini

Architecture, Planning, and Design

Poster abstract:

This research explores how artificial intelligence can interpret human perceptions of urban environments to inform more inclusive, comfortable, and safe design outcomes. Using street view imagery and machine learning models, we analyse spatial qualities through a human-centred lens, translating subjective experiences into data-driven insights for better city-making.

Poster keywords: Perception - AI - Explainable AI (XAI)

Focus area(s): Human-Centred AI;Data-Centric AI

Supervisors: Dr. Silvia Micheli - Dr. Dan Luo - Prof. Yan Liu

Model-Based Transfer Learning for Traditional ML Models in Clinical Prediction Tasks

Oliver Pienaar

CHSR

Poster abstract: Model-based transfer learning (MBTL) enables machine learning (ML) models to adapt to new settings with limited data–especially in healthcare, where privacy and resource constraints often hinder local model development. While established in deep learning, MBTL remains underexplored for traditional ML models widely used in healthcare, such as logistic regression, random forests, and XGBoost. This study evaluates MBTL across these models on two clinical prediction tasks: in-hospital mortality (classification) and remaining length of stay (RLOS; regression). Using the eICU Collaborative Research Database, we simulate hospital-to-hospital transfer under varying source dataset types (homogeneous vs. heterogeneous) and target sizes (50–1400 samples). We compare three strategies: fine-tuning (MBTL), re-training from scratch, and direct reuse ("as-is"). MBTL consistently outperforms the others when the source dataset is larger than the target, particularly for high-capacity models and regression tasks. Random forests showed strong MBTL performance, while logistic regression benefited less from fine-tuning. Unexpectedly, homogeneous source datasets often yielded better transfer performance than similarly sized heterogeneous ones. These findings offer empirical benchmarks and guidance for deploying ML in data-constrained clinical environments, highlighting MBTL as a practical pathway for small hospitals to adopt ML tools without large local datasets.

Poster keywords: Machine Learning, Transfer Learning, Healthcare

Focus area(s): Scalable and Sustainable AI;Data-Centric AI

Supervisors: Jason Pole, Moji Ghadimi, Sally Shrapnel, Anton van der Vegt

Automated 3D segmentation of glomeruli in human kidney tissue specimens using 16.4 T MRI without contrast agents.

Aurel Amar

AIBN, EECS

Presentation abstract: Accurate and fast assessment of glomerular number in ex vivo human kidney specimens is challenging. Glomerular number is traditionally estimated using measurements on glomerular profiles, the two-dimensional samples of glomeruli seen in histological sections. Such an approach relies on complex tissue preparation protocols, specialised equipment, and staff trained to manually count glomeruli. It is labour intensive, time-consuming, and can damage the samples studied, limiting complementary analysis with other methods. In this study, we show that glomeruli can be imaged in human nephrectomy samples at ultra-high field (16.4 T) MRI without contrast agents. A three-dimensional segmentation pipeline was developed to use high-resolution MR images (30 μm isotropic resolution) to accurately estimate the number of nephrons in a sample in a fully automated manner. An independent stereological method, the physical disector/fractionator combination, was used to validate the estimates of glomerular number obtained using MRI. Statistical analyses showed a good agreement between the two methods.

Presentation keywords: Glomerular number, MRI, stereology, volumetric image segmentation

Focus area(s): Human-Centred AI;

Supervisors: David Reutens, Nyoman Kurniawan, Shakes Chandra