UQAI
Research Network

2026 UQ AI HDR Showcase

Showcasing AI Research

Hawken Engineering Building (50-S201)
10-11 June 2026

About

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

On behalf of the organising committee, we wish you an enjoyable event!

- Dr Alina Bialkowski
- A/Prof. Rocky Chen
- Professor Shane Culpepper
- Professor Shazia Sadiq

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

Email: at uqai@uq.edu.au

Program

Day 1 – Wednesday 10 June

Time	Session	Speakers	Page
9:00	**Registration + Tea/Coffee**		
9:30	**Event Opening**		
	Opening address	Prof. Pierre Benckendorff	
	AI Research Network	Prof. Shazia Sadiq	
	Event preview	Dr Alina Bialkowski	
10:00	**Keynote Presentation**	Dr Stefan Hajkowicz	5
11:00	**Short break (no refreshments)**		
11:10	**Presentation Session 1**		
	LLMs as Scalable Qualitative Interviewers	David Cheney	13
	Design and Evaluation of an AI-Mediated Conversational Tool to Elicit Evidence of Student Understanding of Submitted Work	Karen Kim	22
	Audio Retrieval: Challenges and Solutions	Watheq Mansour	36
	Are Vision-Language Models Reliable Concept Annotators? A Neuron-Level Analysis	Saied Salem	33
	VAEye: Mapping the Gap Between Native and In Vitro Tissue Models via Latent Space Manifold Analysis	Diya Choudhary	14
	Hybrid Machine Learning-Population Pharmacokinetic Modelling for Precision Antibiotic Dosing: Lessons from Data Quality and Opportunities for the ROAD RCT	Rob Oakley	31
12:00	**Lunch**		
13:00	**Panel session: How is AI reshaping the way we work and live**	Dr Aninda Saha, Dr Ashleigh Richardson, Dr Paul Vrbik, Dr Shengyao Zhuang	8
14:00	**Afternoon tea break + poster/demo session**		
	Ethics in the Age of AI	Michael Tafe	39
	Enhancing Lexical Resources for Argumentative Essay Writing through Corpus and AI Integration	Nguyen Thi Thuy Linh	40
	SAMRI: Segment Any MRI — A Foundation Model for Whole-Body MRI Segmentation	Zhao Wang	42
	Generating High-Fidelity Semi-Synthetic Concept Drift Streams for Electronic Medical Records	Oliver Pienaar	41
	AI in the detection of preclinical and early osteoarthritis	Jessica Perchman	37
	AI Parent Selection: Sensory and agronomic traits in strawberry (Fragaria x ananassa)	Laura Esquivel Garcia	38
	Audio Retrieval: Challenges and Solutions	Watheq Mansour	36
15:00	**Presentation Session 2**		
	Smart Sowing: Intelligent Agronomic Decision Support for Sustainable Agriculture	Kaveesha Prabhath	23
	Red Lines: Intolerable AI Thresholds Informed by the Global Public and AI Experts	Gerard Hassed	17
	Development of Agentic AI for the Automated Calibration of Soft Sensors	Min-Chia Lan	26
	Ripple Grader: Teacher-Supervised, Rubric Aligned AI Assessments	Nicholas Nucifora	28
	Machine Learning-Based Deformation Forecasting for Mine Waste Infrastructure: An Exploratory Framework for InSAR-Meteorology Integration	Augusto Riascos Realpe	10
	AI-assisted Decision-making in Collaborative XR	Quinn Zhu	30
	Not MY Son: Sympathy For AI (Or, We Love WALL-E)	Talia Parker	34
	From Papers to Knowledge: Intelligent Scholarly Data Management at Scale	Daomin Ji	12
16:00	**Day 1 wrap-up**		

Day 2 – Thursday 11 June

Time	Session	Speakers	
9:20	**Day 2 Opening**	A/Prof. Rocky Chen	
9:30	**Presentation session 3**		
	AI at Work: Gendered Career Trajectories in Banking's Digital Transformation	Naiyer Fatema Khanom	27
	FOSTER: First-order Dataset Distillation for Text-based Sequential Recommendation	Hung Tran	18
	Beyond Translation: Using AI in Fieldwork	Jenny Tang	19
	Surfacing AI Reliance in Real-Time: Supporting Metacognition and Reducing Cognitive Offloading in LLM-Assisted Education	Martin Komitski	25
	Toward a Participatory Approach in the Implementation of AI in Colombian Universities: Analysis of Protocols and Proposals for Knowledge Production	Juan Camilo Ospina Deaza	21
	Data Discovery for Reliable AI over Complex Tabular Data	Feng Luo	15
	Practical AI Adoption in Industry: Applications and Challenges in Everyday Business Use	Brendan Edmonds	11
10:30	**Morning tea**		
11:00	**Keynote – Beyond the Algorithm: Building AI Literacy for a Society That Works for Everyone**	Sophia Arkinstall	6
12:00	**Lunch**		
13:00	**Interactive tutorial: Building your own research apps with vibe coding**	Dr Aneesha Bakharia	7
14:00	**Afternoon tea**		
14:30	**Presentation session 4**		
	Teaching Claude to See 'Like': A Human-in-the-Loop Annotation Story	Aisha Aslam	9
	Benchmarking Training-Free Token Compression for Large Vision Language Models in Agriculture	Rukshan Karannagoda	32
	Catching Prompting Bias Live: How People Use AI During Cross-Ideological Vaccination Debates	Thomas Felesina	35
	Revealing Shared Data: Table Overlap Discovery for Data-Centric AI	Ge Lee	16
	Population Structure in Australian Macadamia Parental Germplasm	Ni-Jung Lin	24
	AI in Personalising Customer Experience	Joel Andrade	20
	Towards Responsibility-Aware Human-Centred Generative AI Design for Older Adults' Health and Wellbeing	Pradeepa Bandara	29
15:30	**Trivia**	Prof. Shane Culpepper	
15:45	**Closing remarks and awards**	Dr. Alina Bialkowski & Rocky Chen	

Keynote Presentations

Becoming a “Purple Person": Rewriting Australia’s AI Roadmap

Dr Stefan Hajkowicz
Chief Research Consultant, CSIRO

Abstract:
When CSIRO developed Australia’s national AI roadmap and ethics framework in 2019, even the experts underestimated how quickly the technology would advance in usability and accuracy. Reflecting on the “spine-tingling" leaps in AI capability over the last year, this presentation addresses what a roadmap “do-over" must include: a critical focus on jobs, skills, and workforce transition.

As our country faces compounding challenges, including climate variability, rising production costs, and severe skills shortages, the key to unlocking true organisational productivity lies in how we adapt. This talk introduces the concept of the “purple person" as the future of the workforce. Drawing from a 10-year CSIRO study of 75,000 job advertisements, this talk explores why the most valuable future professionals will combine “red" technical skills (machine learning, data analysis) with “blue" business acumen (communication, leadership, and innovation). Attendees will discover how to transition from being merely amazed by AI to practically mastering it.

Biography:
Through his research, consulting and advisory work, Stefan helps government, industry and community organisations explore the future and make wise choices. He focuses on supporting data-driven approaches to strategy, foresight and decision making. Much of his work is in the area of translational research; about how knowledge discoveries are translated into real-world solutions and better decisions with better outcomes. In recent times he has been working on artificial intelligence (AI) policy and strategy. He authored the Australian national AI roadmap and ethics framework and has been studying the diffusion of AI across professions and research fields. He is also working on the economic geography of Australia’s rapid-growth digital technology industries and the new digitally enabled geography of work. He loves interdisciplinary research and the most rewarding aspect of his job is seeing the discovery of knowledge result in better quality-of-life for people.

Back to Program

Beyond the Algorithm: Building AI Literacy for a Society That Works for Everyone

Sophia Arkinstall
Founder, AI & Society

Website: https://aiandsociety.io

Abstract:
Artificial intelligence is no longer confined to research labs and technology companies. It is shaping how we learn, work, govern, create, and make decisions. Yet while AI capability is advancing rapidly, public understanding, institutional readiness, and societal trust often struggle to keep pace.

Drawing on her experience founding AI & Society, a national community exploring the societal impacts of artificial intelligence, Sophia Arkinstall will explore what it means to build AI literacy in an age where everyone—not just technologists—will be affected by AI.

For HDR students and emerging researchers, this keynote offers a broader perspective on the role of research beyond publication. How do researchers communicate complex ideas to the public? What responsibilities come with developing powerful technologies? And what skills will be required to bridge the gap between technical innovation and societal impact?

This session will explore the future of AI literacy, human judgement, interdisciplinary collaboration, and the role researchers can play in shaping an AI-enabled society that works for everyone.

Biography:
Sophia Arkinstall is the Founder and Director of AI & Society, a national community exploring the societal impacts of artificial intelligence. Since launching in Brisbane in January 2025, AI & Society has grown to over 3,000 people across five cities — a cross-disciplinary network connecting people from academia, industry, government, and the broader community through events, workshops, and public conversations about AI.

With a background spanning smart cities, innovation, education, health, and sustainability, Sophia is passionate about helping people navigate the social, economic, and human dimensions of technological change. Her work sits at the intersection of AI, society, and place, informed by global study tours and engagement with leading technology, research, and innovation ecosystems around the world.

Sophia speaks on AI literacy, future skills, responsible AI adoption, and the role of communities in shaping AI futures.

Interactive Tutorial

From Research Idea to Web App: Responsible Prototyping with Agentic AI

Aneesha Bakharia
Senior Lecturer, UQ

Abstract:
Agentic AI tools are making it easier for researchers to turn ideas into interactive web applications, explainers, dashboards, and prototypes. These tools can support the development and distribution of research-facing applications, but they are most useful when guided by clear specifications, testing practices, and responsible human oversight.

In this interactive session, participants will learn how to use agentic AI for research app development. The session will introduce practical approaches to spec-driven development, describing backend logic, and using AI tools to generate, test, and refine a working prototype. The session will also cover strategies for reviewing AI-generated code and keeping the researcher in control of important design decisions.

This is an interactive tutorial. Participants are encouraged to bring a **laptop** with access to an **agentic AI coding tool**, such as Claude Code or OpenAI Codex, or a general-purpose AI chatbot. For those without access, an open-source option will be provided.

Biography:
Dr Aneesha Bakharia is a Senior Lecturer in the School of Electrical Engineering and Computer Science at The University of Queensland. Her background spans electronics engineering, educational software development, learning analytics, and technology-enhanced learning. She has written nine books on programming and web development, and her current work focuses on generative AI in education, AI-assisted programming, learning analytics, and responsible AI use in teaching and assessment.

Back to Program

Panel Discussion

Panel Session: How will AI Reshape the Way We Work and Live?

***Dr Aninda Saha*[1], *Dr Ashleigh Richardson*[2], *Dr Paul Vrbik*[3], *Dr Shengyao Zhuang*[4]**




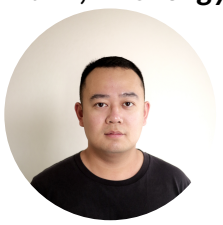

1: Senior AI Engineer, Canva Research 2: Director & Software Engineer, Persephia
3: Senior Lecturer, School of EECS 4: Founding Member, Silicon Valley startup

Aninda Saha is a Senior AI Researcher at Canva, where he works on generative AI and creative intelligence systems, and contributed to the Canva Design Model and Magic Layers. He completed his PhD in Computer Vision at UQ in collaboration with CSIRO's Data61, focussing on efficient deep learning, model compression and knowledge distillation. Prior to joining Canva, Aninda was a Senior AI Researcher at Leonardo.AI, where he played a key role in developing Phoenix, Australia's first foundational image model. His work spans computer vision, multimodal AI and generative models.

Ashleigh Richardson is Director of Persephia, and holds a Bachelor of Software Engineering and a PhD in Artificial Intelligence from UQ. Her research is centered around addressing challenges that have arisen from the recent rapid developments in the field of natural language processing. Ashleigh has a strong passion for teaching, and in addition to her role at Persephia Ashleigh lectures computer science at UQ. Aside from her work in machine learning and education, Ashleigh is currently focused on projects in the areas of web and game development.

Paul Vrbik is a Senior Lecturer in the School of EECS at UQ. His teaching and research span software engineering, computer science education, and computational mathematics. Dr Vrbik is known for his innovative teaching approach, particularly his use of live-coding demonstrations to develop students' problem-solving and software development skills, earning him a University of Queensland Award for Teaching Excellence. His research interests include computing education, academic integrity in computer science, algorithm design, and mathematical computing.

Shengyao Zhuang is a founding member of a Silicon Valley startup building agentic search systems for financial markets. Previously, he was an Applied Scientist in Amazon's AGI labs, where he contributed to web search technologies that support Amazon's LLM-powered products. Before joining Amazon, he was a Postdoc at CSIRO's Australian e-Health Research Centre, developing LLM-driven search systems for healthcare applications. He received his Ph.D. in Computer Science from UQ, under the supervision of Prof. Guido Zuccon. His research interests span information retrieval, large language models, AI agents, neural ranking, and natural language processing.

List of Presentations

Teaching Claude to See 'Like': A Human-in-the-Loop Annotation Story

Aisha Aslam

School of Languages and Cultures

Contact: https://www.linkedin.com/in/aisha-aslam/

Presentation abstract: Certain lexico-grammatical features resist automated annotation due to pragmatic complexity, creating gaps in corpus analyses such as Multi-Dimensional Analysis (MDA). The word like is among the most challenging, realising functionally distinct categories including figurative simile, literal comparison, exemplification, preference, and discourse marking. Yet existing taggers classify it as a 'bin tag', excluding it entirely from statistical output. This research reports a human-AI collaborative annotation protocol using the Three-Minute Thesis Corpus (3MTC; 454321 total tokens, 1,619 LIKE tokens) as a test case. A large language model (Claude, Anthropic) produced a first-pass annotation across an initial four-category taxonomy. Human verification rounds identified systematic error patterns and recalibrated the model's decision rules. Crucially, human correction revealed a fifth category systematically misclassified by the AI, demonstrating that human oversight can reshape the taxonomy itself, not merely improve accuracy within a fixed set. Inter-rater reliability improved from moderate to substantial agreement (κ = 0.549 to κ = 0.776). Distributional analysis shows exemplification dominates at 50%, consistent with the 3MT as a science popularisation genre. For the AI HDR community, this protocol offers a replicable, low-cost solution for any MDA study where theoretically motivated features fall outside existing tagset boundaries.

Presentation keywords: Human-AI collaboration, corpus linguistics, Multidimensional Analysis

Focus area(s): Human-Centred AI;

Supervisors: As/Pr Peter Crosthwaite, Dr Martin Schweinberger

Machine Learning-Based Deformation Forecasting for Mine Waste Infrastructure: An Exploratory Framework for InSAR-Meteorology Integration

Augusto Riascos

Civil Engineering

Contact: www.linkedin.com/in/augustoriascos

Presentation abstract: Subsidence monitoring in tailings storage facilities and waste rock dumps requires improved early-warning capability beyond traditional reactive approaches limited by satellite acquisition intervals. We present an exploratory machine learning framework combining Synthetic Aperture Radar interferometry time series with meteorological data to investigate feasibility of multi-day deformation forecasting. Our methodology employs temporal feature engineering through lag decomposition and derivative signal processing, evaluated across six regression algorithms using walk-forward cross-validation on operational mining sites. Natural deformation variability establishes baseline performance reference. Results reveal significant site-specific variations in predictive skill, with persistent scatterers consistently outperforming distributed scatterers. Predictive performance ranges from marginal improvements to substantial degradation relative to baseline, with feature optimization emerging as critical site-dependent factor. We identify key meteorological and InSAR-derived drivers of deformation and document constraints limiting generalization. Findings suggest meteorological integration holds promise for enhanced forecasting but requires feature customization and validation before operational deployment. This work establishes methodology for investigating site-specific deformation mechanisms and supports evidence-based refinement of monitoring strategies for mine waste risk assessment.

Presentation keywords: InSAR forecasting, Mine waste monitoring, ML deformation prediction, Meteorological integration, Subsidence prediction

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

Supervisors: Prof. David J Williams

Practical AI Adoption in Industry: Applications and Challenges in Everyday Business Use

Brendan Edmonds

EECS

Presentation abstract: The increasing accessibility of AI technologies has accelerated their adoption across a broad range of industries. Organisations are integrating AI into daily workflows for tasks including automation, analytics, software engineering, content generation, and operational support. Despite these advances, significant challenges remain in achieving reliable and effective use within real-world business environments. This presentation investigates current industry applications of AI and analyses practical limitations associated with its use, including model reliability, explainability, governance, privacy, human factors, and organisational integration. The discussion focuses on lessons learned from real-world usage and considers approaches for balancing innovation with responsible and sustainable AI adoption.

Presentation keywords: Business Integration, AI Governance, Practical Challenges

Focus area(s): AI Strategy and Governance;

Supervisors: Mark Utting

From Papers to Knowledge: Intelligent Scholarly Data Management at Scale

Daomin Ji

EECS

Contact: https://daominji.github.io/homepage/

Presentation abstract: The modern scholarly record now spans millions of papers across rapidly evolving disciplines, yet much of this knowledge remains difficult to organize, search, and act upon. This presentation introduces our work on intelligent scholarly data management, which seeks to transform large-scale paper repositories into structured, navigable, and actionable research infrastructure. We approach the problem from three complementary perspectives. First, we investigate how to organize large paper collections into topic-aware structures that reveal the conceptual landscape of a research area. Second, we develop methods to support scholarly service tasks — such as paper–reviewer assignment and expert discovery — by jointly modeling the relationships among papers, topics, and researchers. Third, we build cross-disciplinary scholarly knowledge infrastructure that enables research agents to answer complex queries and to reason over the broader scientific literature. Together, these directions move beyond traditional paper search toward intelligent systems that can structure, connect, and reason over scholarly knowledge at scale.

Presentation keywords: Scholarly Data, Taxonomy, Research Land Scape

Focus area(s): Data-Centric AI;

Supervisors: Prof. Zhifeng Bao

LLMs as scalable qualitative interviewers

David Cheney

Business School

Presentation abstract: Beliefs towards climate change are a topic of interest in social science, yet little research has examined why people change their minds about climate change. This project integrates a custom-built web application with the OpenAI API to function as a qualitative interviewer. Five hundred and one (N = 501) interviews were collected. Participants self-classified into one of seven belief-trajectory categories (e.g., climate change does not exist → does exist). Thematic analysis reveals several distinct themes among current believers: Observation of weather (e.g., increasing extreme weather); Exposure to new information (e.g., scientific evidence); Observation of human impact (e.g., causal links between pollution and climate); and Social influence (e.g., conversations with friends). Three distinct themes are emerging among current sceptics: Hidden agenda (e.g., participants view climate change as a vehicle for political, financial, or elite control); Evidence and education (e.g., participants report conducting their own research or concluding that predicted disasters failed to materialise); and Social influence (participants are influenced by friends, family, or dislike of climate advocates). This study produces one of the largest qualitative datasets on climate belief change to date and demonstrates that LLMs can function as scalable and reliably candid qualitative interviewers.

Presentation keywords: Qualitative Research Interview Methods AI-Interviewer

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

Supervisors: Prof. Matthew Hornsey, Dr Samuel Pearson

VAEye: Mapping the Gap Between Native and In Vitro Tissue Models via Latent Space Manifold Analysis

Diya Choudhary

SCMB

Contact: https://www.linkedin.com/in/diya-prabhuram-61a30a197

Presentation abstract: Retinoblastoma (RB) cell lines are widely used as experimental models of retinal biology, yet their transcriptomic fidelity relative to native tissue remains poorly characterised. Here we present VAEye, a variational autoencoder trained on 1028 ocular transcriptomic profiles from the National Eye Institute data portal, encompassing native retina, native RPE, and cell line-derived samples. Without explicit supervision, VAEye learned a compact 16-dimensional latent representation in which native and cell line samples occupy distinct regions, with four active latent dimensions identified by KL divergence thresholding capturing the dominant axes of biological variation. To interpret what these dimensions encode, we applied Integrated Gradients through an encoder-MLP pipeline, attributing the model's native-versus-cell-line discrimination to individual genes and connecting attributions to differential expression analysis and gene set enrichment. Projection of two DepMap retinoblastoma cell lines into the trained latent space revealed that RB lines are near-orthogonal to native retinal tissue. Decoder-based gene displacement further characterised the transcriptional gap, identifying loss of photoreceptor identity and gain of proliferative and metabolic signatures in RB lines. This work demonstrates that explainable deep learning applied to transcriptomic data can quantify biological drift between tissue states, offering a principled framework for cell line fidelity assessment.

Presentation keywords: XAI, Transcriptomics, VAE, Deep Learning

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

Supervisors: Prof. David Ascher

Data Discovery for Reliable AI over Complex Tabular Data

Feng Luo

EECS

Contact: https://scholar.google.com/citations?user=OH8VnF8AAAAJ&hl=zh-CN&oi=ao

Presentation abstract: Real-world tabular data is rarely clean, complete, or ready for direct analysis. It is often incomplete, scattered across large table collections, and represented in complex formats. To reliably support analytical tasks, AI systems must first discover the right evidence, structure, and context from such data. This motivates my research topic: data discovery for reliable AI over complex tabular data.

The key challenge is that useful tabular information is often implicit, fragmented, or hidden in heterogeneous data environments. As a result, AI systems may rely on incomplete evidence, miss important structural context, or reason over unreliable information.

My research addresses this challenge by developing methods that make tabular evidence more discoverable and usable for reliable AI analytics. Specifically, I study evidence discovery under incomplete data, multi-table discovery over large table collections, and structure discovery for semi-structured tables. Together, these works aim to move AI systems beyond analyzing prepared tables, towards reliably discovering and using evidence from complex real-world tabular data.

Presentation keywords: Data Discovery; Reliable AI; Complex Tabular Data; Tabular Analytics

Focus area(s): Data-Centric AI;

Supervisors: Zhifeng Bao, J. Shane Culpepper

Revealing Shared Data: Table Overlap Discovery for Data-Centric AI

Ge Lee

EECS

Contact: https://www.linkedin.com/in/ge-lee/;https://gelee-cs.github.io/

Presentation abstract: Modern AI systems need more than better models. They need data that can be trusted, traced, and reused. In large repositories, however, tabular data are often fragmented across tables that have been copied, revised, partially reused, or reorganised over time. As a result, two tables may contain substantially overlapping information while appearing unrelated because their layouts differ, metadata are missing, or only part of the content is shared. To address these challenges, we make shared tabular content easier to detect, even when tables are not aligned. This enables systems to uncover hidden relationships between tables and rapidly identify relevant tables at scale before expensive downstream verification. The outcome is a practical foundation for data-centric AI, making repositories easier to maintain, reused data easier to trace, redundant content easier to find, and evolving datasets easier to compare, while supporting more reliable data curation at lower computational cost.

Presentation keywords: Table overlap; Scalable data curation

Focus area(s): Data-Centric AI;

Supervisors: Zhifeng Bao

Red Lines: Intolerable AI Thresholds Informed by the Global Public and AI Experts

Gerard Hassed

School of Psychology

Contact: https://www.linkedin.com/in/gerardhassed/

Presentation abstract: AI systems are becoming more capable and dangerous. Many experts agree that there is at least some risk of catastrophic harms from AI. To prevent this, there is a global call for the establishment for international AI red lines: the uses and capabilities of AI that should be banned outright. While society is determining the prohibited uses of AI, one aim of this thesis is to get on these 'capability' red lines. Drawing a red line on AI capabilities is considerably more difficult. An international scientific panel has proposed five core red lines: Autonomous Replication or Improvement; Power Seeking; Assisting Weapon Development; Cyberattacks; Deception. However, the definition and scope of these red lines remain contested. What constitutes an intolerable level of cyberattack skill or deception? Where should the line be drawn? How can red lines be enforced? Through public and expert consultation, as well as a gap analysis on current prohibitive governance, my thesis seeks to establish what the AI red lines should be, how they can be measured, and how they can be enforced.

Presentation keywords: Regulation, Governance, Safety

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

Supervisors: Associate Professor Michael Noetel, Professor Nicole Gillespie, and Dr Steven Lockey

FOSTER: First-order Dataset Distillation for Text-based Sequential Recommendation

Hung Tran

EECS

Contact: https://chenxing1999.github.io/

Presentation abstract: Text-based sequential recommender systems, while achieving strong performance, are undeniably more expensive to train. Dataset distillation can alleviate this bottleneck by condensing the large dataset into a compact set of synthetic samples. Yet, applying dataset distillation to text-based sequential recommendation is not trivial due to high computational cost of the language models and the more complex autoregressive objective compared to standard classification. To this end, we propose **FOSTER** – **F**irst-**o**rder data**s**et distillation for **Te**xt-based Sequential **R**ecommendation – addresses these challenges via: (1) stochastic item subset sampling, replacing costly full-corpus embedding extraction; (2) a first-order optimisation with trajectory-anchored parameter reset to avoid expensive bi-level gradients; and (3) a novel regularization loss that explicitly encourages the distributional hypothesis. Extensive experiments show that FOSTER consistently outperforms baselines, approximating full-dataset performance using only 20 synthetic sequences for Games and Foods, and 60 for Yelp-Nashville.

Presentation keywords: Recommendation, Dataset Distillation, Training Efficiency

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

Supervisors: A/Prof Rocky Chen

Beyond translation: Using AI in fieldwork

Jenny Tang

School of Social Science

Contact: https://www.linkedin.com/in/jennytang58/

Presentation abstract: This talk offers a personal and recent account of the practical usages of AI in the planning and the use during a multilingual fieldwork. AI tools are increasingly present in research workflows, but how useful are they when your data exists in multiple languages, culturally specific applications, and partially untranslatable concepts? This talk draws on 4 months of ethnographic fieldwork in Yunnan, China, to share anecdotes and examples of how AI reshaped the research process and where it fell short. First, AI compressed the timeline for producing bilingual ethics and participant materials, but required significant human judgement to catch mistranslations that were technically correct but culturally wrong. Second, AI became a useful thinking partner for working through qualitative data - not as an analytical shortcut, but as a sounding board for interpretation across languages. But most interestingly, AI struggled with foreign concepts - exposing the limits of tools trained predominantly on English-language academic conventions.

Presentation keywords: Anthropology, fieldwork, cultural

Focus area(s): Human-Centred AI

Supervisors: Diana Young

AI in Personalising Customer Experience

Joel Andrade

Business School

Contact: www.linkedin.com/in/joel-andrade-thought4food

Presentation abstract: Research on personalisation, though still limited, spans multiple disciplines and highlights its potential to enhance customer value and firm profitability. Industry surveys similarly suggest that nearly three in four consumers are more likely to purchase from brands delivering personalised experiences. Despite this, firms continue to struggle with achieving return on investment from personalisation efforts. The rise of Artificial Intelligence (AI), particularly Generative AI such as chatbots and Large Language Models (LLMs), is expected to transform customer experience (CX) through more adaptive and personalised interactions. However, existing service research presents fragmented findings, indicating a limited understanding of how AI-driven personalisation is applied within service contexts. Responding to Ostrom et al.'s call for frictionless digital experiences, this study explores the intersection of AI, personalisation, and CX through an integrative review of 88 articles reflecting consumer perceptions of AI-driven personalisation. The review synthesises insights across key service research themes, including customer engagement, trust and transparency, customer wellbeing, and the privacy–personalisation paradox. Building on these findings, the study proposes a novel typology for AI-personalised CX alongside the ENACT roadmap, designed to help managers strategically leverage AI for personalisation, along with a research agenda for future research.

Presentation keywords: AI, Personalisation, Customer Experience

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

Supervisors: Professor Janet McColl-Kennedy, Associate Professor Len Coote

Toward a participatory approach in the implementation of AI in Colombian Universities: analysis of protocols and proposals for knowledge production

Juan Camilo Ospina Deaza

School of Communication and Arts

Contact: https://www.linkedin.com/in/juan-ospina-91581525a/

Presentation abstract: This article examines the development and implementation of artificial intelligence (AI) protocols in Colombian universities, within a context marked by the digital divide and educational challenges. Using a grounded theory approach, eleven institutional documents published between 2023 and 2025 were analyzed, identifying key themes such as ethical responsibility, transparency, data protection, and equitable access. The protocols reveal divergences in the definition of AI, ranging from functional to critical and humanistic approaches. However, significant omissions are observed, including the lack of risk assessment frameworks and the exclusion of diverse actors from the educational community. The creation of AI protocols in academic contexts implies rethinking how knowledge is constructed, validated, and shared. While AI can transform teaching, research, and management processes, it also poses risks such as the dehumanization of learning, the perpetuation of algorithmic biases, and the exclusion of marginalized communities. This study proposes a roadmap for updating and creating new protocols, integrating participatory and interdisciplinary approaches to ensure the inclusion of diverse voices and a comprehensive understanding of the AI lifecycle, beyond the use of generative models.

Presentation keywords: Colombia, IA, participation, cocreation

Focus area(s): AI Strategy and Governance; Scalable and Sustainable AI;

Supervisors: Giselle Newton

Design and Evaluation of an AI-Mediated Conversational Tool to Elicit Evidence of Student Understanding of Submitted Work

Karen Kim

School of Education

Contact: https://www.linkedin.com/in/karenglish2017/

Presentation abstract: Generative artificial intelligent (GenAI) has reduced the reliability of written submissions as indicators of student understanding due to its ability to generate human-like texts (Corbin et al., 2025). Oral interviews are considered a promising mechanism as they require students to explain the reasoning and processes behind their work (Ward et al., 2024). However, they are difficult to implement at scale (O'Riordan et al., 2025). Therefore, this doctoral research aims to design and evaluate an AI-mediated oral interview tool conducted post-submission in open assessment contexts. Using a two-phase design-evaluate methodology, this study conceptualises the assessment design based on the Evidence-Centred Assessment Design principle (Mislevy et al., 2017) and operationalises it through Finite State Machine-based AI prompting to generate personalised oral questions from each student's submission. In the design phase, educators and students (N=14) trial the prototype and provide feedback on the question flow and report card logic during interviews, later analysed using Thematic Analysis for system refinement (Braun & Clarke, 2020). In the evaluation phase, this refined system is examined from two perspectives: assessment function and system usability. The study is expected to contribute a theoretically grounded dialogic assessment approach that supports scalable implementation across disciplines.

Presentation keywords: Evidence-Centred Assessment Design, Human-AI Shared Regulation, Conversational AI, Student Understanding

Focus area(s): Human-Centred AI;

Supervisors: A/Prof Hassan Khosravi, A/Prof Christine Slade, A/Prof Peter Crosthwaite

Smart Sowing: Intelligent Agronomic Decision Support for Sustainable Agriculture

Kaveesha Prabhath

QAAFI

Contact: www.linkedin.com/in/kprabhath

Presentation abstract: Agronomic decisions rely on distributed knowledge across trial reports and sowing guides where the meaning of a yield, disease score, or variety trait is tied to a specific crop section, region, and reporting year. Retrieval Augmented Generation (RAG) enables natural-language access to this evidence, but only if the document structure that gives agronomic values their meaning is preserved. When document parsing and chunking separate these dependencies, LLM-based decision support fails due to context mismatch and hallucinated numeric claims. This study examines improved methods for document segmentation and retrieval to preserve agronomic context through layout-aware semantic tagging, dual table representations, metadata filtering, reranking, and an adaptive retrieval management graph–vector hybrid. Semantic tagging preserved context but reduced retrieval separability in a semantically homogeneous corpus. However, metadata-conditioned retrieval optimization significantly improved answer precision from 34.4% in the baseline vector RAG to 82.6% across strata. Elbow-based adaptive retrieval was unstable due to gradual similarity decay, and graph augmentation added latency without yielding consistent gains. Reliable RAG for agronomic decision support depends on structure-preserving ingestion and precision retrieval. Future work will utilize agentic workflows paired with predictive modelling to strengthen reasoning capabilities.

Presentation keywords: Large Language Models, Retrieval Augmented Generation, Agronomic Decision Support, Agricultural Information Retrieval, Sustainability

Focus area(s): Human-Centred AI;

Supervisors: Ass/Prof Sudhir Yadav, Ass/Prof Mahsa Baktashmotlagh, Dr Bruno Rafael, Dr Sebastian Lopez

Population Structure in Australian Macadamia Parental Germplasm

Lin, Ni-Jung (Judy)

QAAFI

Contact: www.linkedin.com/in/linnijung

Presentation abstract: Macadamia, an Australian native subtropical perennial tree crop, presents critical breeding challenges due to inherent heterogeneity, long juvenility, single plant basis progeny evaluation, variation in field trial sites and the complexity of genotype by environment interactions. Aiming to develop and validate an efficient parent selection framework by integrating advanced statistical methods, genomic tools and AI. It delivers comprehensive genomic analysis across all four species of Macadamia: M. integrifolia, M. tetraphylla, M. ternifolia, and M. jansenii. Collected DNA from 2,010 accessions were genotyped by 12,655 DArT markers, revealing a high 45.45% missing markers rate potential driven by deletion, insertion, and allele drop-out.

Informative markers were filtered by call rate, remaining 1,931 markers at 90%. Missing markers were imputed using three algorithms, (1) mode, (2) k-nearest neighbors (KNNi), and (3) random forest (missForest), implemented in R. Imputation accuracy was assessed by randomly masking 5% of the data, 78.60%, 87.14%, and 84.73%, respectively.

Population structure analysis using NbClust in R explained 17.45% to 20.05% of principal component variance. These findings highlight the potential of marker imputation and AI-driven approaches to accelerate genetic improvement, shorten breeding cycles, and enhance genetic gain in macadamia breeding programs.

Presentation keywords: Macadamia, marker imputation, genetic diversity, population structure, artificial intelligence

Focus area(s): Data-Centric AI;

Supervisors: Dr. Mobashwer Alam

Surfacing AI Reliance in Real-Time: Supporting Metacognition and Reducing Cognitive Offloading in LLM-Assisted Education

Martin Komitski

School of Psychology

Presentation abstract: This project introduces a real-time metric for AI reliance, designed to develop more independently capable learners in LLM-assisted educational settings. Despite rapid adoption of LLMs in education, short-term performance gains fail to translate into deep, durable, and transferable knowledge. While considerable effort has been devoted to aligning LLM behaviour with pedagogical principles, little attention has been given to the psychological variables that underpin learner interaction with these tools. A key psychological driver is cognitive offloading; learners delegate knowledge-building to the LLM rather than effortfully engaging with material themselves. Compounding this, the fluency and confidence of LLM responses can distort learners' judgements of their own understanding, undermining accurate meta-representation of their epistemic state. By surfacing the degree of AI reliance in real-time to both the learner and the LLM, this project aims to support more accurate self-assessment and reduce maladaptive cognitive offloading to AI for deep knowledge-building tasks. If successful, this work offers a scalable approach to preserving genuine knowledge construction in an era of increasingly capable AI tools.

Presentation keywords: Learning, Metacognition, AI Reliance

Focus area(s): Human-Centred AI;

Supervisors: Dr David Sewell, Dr Hassan Khosravi

Development of Agentic AI for the Automated Calibration of Soft Sensors

Min-Chia Lan

EECS, SMI

Contact: www.linkedin.com/in/melody-lan-findingjobs

Presentation abstract: SAG (Semi-Autogenous Grinding) mills are critical to mineral processing, but key operational parameters, such as ball filling and total filling, cannot be measured directly during continuous operation. The Mill Filling Inference Tool (MillFIT), developed by JKMRC at The University of Queensland, estimates these quantities from existing instrumentation. However, deploying MillFIT at a new site remained labor-intensive, requiring manual extraction of tables from liner scan PDFs and Excel reports, cleaning of historical exports, and transcription into calibration templates. This project implemented an agentic AI system to automate this data preparation workflow. Built on Google's Agent Development Kit (ADK), the system used an orchestrator–subagent architecture in which a root agent assigns tasks to specialized sub-agents for file ingestion, table extraction, and cloud storage export. Users interacted via a conversational chatbot to upload documents or provide cloud storage references, and extracted tabular data was presented for confirmation before being written to a designated GCS bucket. The system also supported web search, report generation, and feedback capture, with session persistence and browser-based deployment on Google Cloud Platform.

Presentation keywords: Agentic AI, Mineral Processing, Google ADK

Focus area(s): Data-Centric AI;

Supervisors: Dr. Christian Zuluaga Bedoya

AI at Work: Gendered Career Trajectories in Banking's Digital Transformation

Naiyer Fatema Khanom

School of Social Science

Contact: https://www.linkedin.com/in/naiyerkhanom/

Presentation abstract: AI is rapidly reshaping the future of work in banking—redefining roles, skills, and career pathways. This paper draws on my PhD research on women's career trajectories in a digitally transforming banking sector in a developing country (Bangladesh), based on 31 semi-structured interviews with professionals across different institutional contexts. While the research focuses on career progression, a key finding is the growing centrality of AI in shaping employees' career outlooks. Participants consistently frame AI as both an opportunity and a risk. AI-driven transformation is generating demand for new digital capabilities, making AI literacy and continuous upskilling critical for career advancement. At the same time, increasing automation of routine roles is raising concerns about job displacement, role reduction, and uncertain advancement pathways. Importantly, the capacity to adapt to AI-enabled change is uneven. Women often face greater constraints in engaging with AI-related learning due to limited time, sociocultural expectations, and restricted access to informal networks. This creates a critical gap between access to digital transformation and the ability to translate it into career progression. By conceptualising AI as an embedded organisational force, this research highlights the need for inclusive and equitable strategies, showing that without targeted support, AI-driven transformation risks reinforcing existing gender inequalities in the future of work.

Presentation keywords: AI-Driven Change, Reskilling & Upskilling, Fintech Transformation

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

Supervisors: Dr Zoe Staines and Dr Lutfun Nahar Lata

Ripple Grader: Teacher-Supervised, Rubric Aligned AI Assessments

Nicholas Nucifora

EECS, Business School

Contact: nicholasnucifora.com

Presentation abstract: Ripple Grader is an AI-assisted assessment platform designed to help educators grade large volumes of RiPPLE student resources and moderations, while preserving teacher control. The system converts rubrics into a structured format that is easier for AI to apply consistently, supports rubric-specific marking notes, allows optional topic-specific context, and lets teachers tune overall marking instructions and the student-facing feedback format.

A key design principle is supervised automation: teachers can preview AI marking on small samples, inspect grade distributions, refine their prompts or rubrics, and only approve full-batch grading once they are satisfied with the output. The platform additionally supports separate rubrics for resources and moderations, cutoff dates and extension handling, custom grade outputs, manual teacher overrides, re-running grade runs, and student-ready feedback emails.

This project explores how AI can reduce repetitive assessment workload, while keeping educator judgement central. It also opens possibilities for new types of assignment designs that may have previously been impractical because grading was too labour-intensive. An initial pilot was run in one course with positive feedback, and the next stage will extend this to four courses and approximately 1000 students. A continuing priority is making AI use transparent to students, showing how it supports learning, while keeping teachers open to student feedback about the implementation.

Presentation keywords: AI Assessment, AI Rubric Grading

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

Supervisors: Russell Manfield

Towards Responsibility-Aware Human-Centred Generative AI Design for Older Adults' Health and Wellbeing

Pradeepa Bandara

EECS

Contact: https://www.linkedin.com/in/pradeepa-bandara-22114722

Presentation abstract: Generative Artificial Intelligence (GenAI) technologies are increasingly being integrated into health, wellbeing, and self-care contexts. However, many AI systems are not designed with older adults' needs, values, and lived experiences in mind. This research explores how older adults imagine, evaluate, and co-design future GenAI technologies to support everyday health and wellbeing. Using a Human-Computer Interaction (HCI) and participatory design approach, the study involves co-design workshops with older adults to investigate opportunities, concerns, and design expectations surrounding GenAI in later life. Preliminary findings indicate that older adults see potential value in GenAI systems that simplify health information, support memory and organisation, assist with medical appointments, encourage healthy routines, and provide personalised wellbeing support. Participants also highlighted concerns relating to trust, autonomy, accessibility, privacy, reliability, and the importance of maintaining human oversight and social connection. The study identifies emerging design considerations for Human-Centred AI systems for ageing populations, including transparency, inclusive interaction design, emotional sensitivity, contextual personalisation, and responsible data practices. This research contributes towards the development of a responsibility-aware design framework for GenAI technologies supporting older adults' health, wellbeing, and self-care.

Presentation keywords: Older Adults, Human Centered AI, Generative AI, Health and Wellbeing

Focus area(s): Human-Centred AI;

Supervisors: A/Prof. Dhaval Vyas, Prof. Tim Miller

AI-assisted decision-making in collaborative XR

Quinn Zhu

EECS

Presentation abstract: As Artificial Intelligence (AI) becomes increasingly embedded in complex, high-risk operational environments, such as surgery, search and rescue (SAR) missions, researchers have typically examined how these systems can generate automated solutions to maximize efficiency. However, these paradigms often regard AI as a replacement for human collaboration, overlooking the importance of shared awareness and decision-making required in collaborative teamwork.

While collaborative Extended Reality (XR) is widely used for dynamic tasks like remote training and maintenance, its application in time-critical SAR operations presents unique challenges in synchronizing perspectives between strategist and pilot. In these settings, effective data visualization is crucial for collaborative decision-making, as it externalizes complex spatial data, establishes common ground, and cultivates shared situational awareness across diverse user interfaces.

This PhD project lies in the intersection of collaborative XR, human-AI interaction, and data visualization, focusing on how to augment human collaborative spatial understanding, decision-making, and performance. The project's ultimate goal is to enhance cognitive coordination and operational success without removing human agency.

Presentation keywords: AI-assisted collaborative distributed interaction

Focus area(s): Human-Centred AI;

Supervisors: Maxime Cordeil, Jarrod Knibbe, Chelsea Dobbins

Hybrid Machine Learning–Population Pharmacokinetic Modelling for Precision Antibiotic Dosing: Lessons from Data Quality and Opportunities for the ROAD RCT

Robert Oakley

School of Medicine

Contact: https://orcid.org/0000-0001-6329-8658,https://www.linkedin.com/in/robert-oakley-32229b6a/

Presentation abstract: Accurate pharmacokinetic (PK) data is essential for antibiotic dosing in critical illness, yet routine clinical datasets contain substantial recording variability that can distort PK parameter estimation and decision-support. My MRes work quantified this issue by comparing vancomycin infusion pump data with electronic prescribing (ePMA) records, showing that administration-time discrepancies changed structural model selection, objective function values, and target breakpoint attainment. My PhD now focuses on a 610-patient international randomised controlled trial evaluating resistance-optimised, Bayesian-guided beta-lactam dosing. Although not an AI trial, it will generate rich PK, microbiology, and clinical outcome datasets suitable for hybrid machine learning (ML)/population PK (PopPK) approaches. Potential applications include ML-driven phenotype discovery, prediction of time-varying clearance, automated PopPK submodel classification, and internet-of-things-enabled real-time data capture to enhance PK model predictive performance. This showcase will present (1) empirical evidence of how data quality shapes PK modelling, (2) conceptual pathways for integrating ML with mechanistic PopPK models, and (3) opportunities to build scalable datasets through collaboration with UQ's SMART Health Data group. My goal is to gather methodological feedback, identify suitable ML techniques, and explore partnerships to advance hybrid PK/AI dosing research.

Presentation keywords: Population Pharmacokinetics, Machine Learning, Data Quality, Precision Dosing, Critical Care

Focus area(s): Data-Centric AI;

Supervisors: Professor Jason Roberts, Dr Aaron Heffernan

Benchmarking Training-Free Token Compression for Large Vision Language Models in Agriculture

Rukshan Karannagoda

EECS

Presentation abstract: Large Vision Language Models (LVLMs) show strong potential for agricultural image analysis and plant phenotyping. However, their deployment in agricultural field environments remains challenging due to the high computational and memory demands of processing high-resolution imagery on edge devices. Existing efficiency techniques do not fully address visual token redundancy at inference time, nor do they examine whether reduction of visual tokens in agriculture is beneficial or harmful. This paper investigates token compression as a practical pathway for improving the efficiency of VLMs in agricultural phenotyping. We review and analyze token reduction strategies for vision-language models and their suitability for high-resolution plant imagery, and show significant accuracy and efficiency increases using token compression along with tiling and augmentation, against three phenotyping tasks. By examining how visual tokens are selected, this work highlights key opportunities and limitations for efficient agricultural VLM inference.

Presentation keywords: Large Vision Language Models; Token Compression; Plant Phenotyping; Precision Agriculture

Focus area(s): Data-Centric AI;

Supervisors: Mahsa Baktashmotlagh, Yadan Luo, Zijian Wang, Scott C. Chapman

Are Vision-Language Models Reliable Concept Annotators? A Neuron-Level Analysis

Saied Salem

EECS

Contact: https://www.linkedin.com/in/saied-salem/

Presentation abstract: Recent Explainable AI (XAI) methods use large language models (LLMs) and vision-language models (VLMs) to generate human-interpretable concepts for explaining image features. However, these explanations often rely on broad contextual cues that may include background noise or misleading patterns. Here, we study whether specific neurons capture concept information. Using a method called Second-Order Lens (SOL), we isolate neural signals linked to defined concepts and use them to segment image regions. Our results show that neurons activations do not always correspond to the correct concept regions, suggesting that current vision-language methods may not be reliable as label-free concept annotators.

Presentation keywords: XAI, Concept-based Explanations, VLM, Mechanistic Interpretability.

Focus area(s): Human-Centred AI

Supervisors: Professor. Tim Miller, Dr. Alina Bialkowski

Not MY Son: Sympathy For AI (Or, We Love WALL-E)

Talia Parker

School of Communication and Arts

Contact: www.linkedin.com/in/talia-parker

Presentation abstract: Portrayals of artificial intelligence in fiction, particularly science fiction films, have real-world consequences for our actual perceptions of these technologies. These portrayals can inspire inventors, cause fear or hope in their audiences, and help us imagine possible futures. One such imagined future takes place in the film WALL-E (2008), where the titular adorable robot develops sentience, falls in love, and ultimately helps the human race restore Earth from its dystopian ruins. This presentation asks how the viewers of WALL-E are brought to feel a sympathetic connection with its main character, and what the formation of this sympathy can reveal about our feelings towards real-world AI. Textual analysis of WALL-E reveals that viewers form a sympathetic connection with him because of his cuteness, his human-like emotions, and his internal development of agency. WALL-E's character presentation can then be compared to the recent disastrous example of a Billie-Jean-dancing robot that became a worldwide laughing stock. Comparing two very different cultural responses to two different types of AI – fictional and real, loved and humiliated – suggests future avenues for examining how people engage with AI in their daily lives. Ultimately, this research calls for continued attention to people's emotional reactions to AI, alongside the development of the technology, to provide a balanced and practical picture of this technology's impact and ongoing integration.

Presentation keywords: Emotion artificial intelligence, sympathy, affective computing, agency, textual analysis

Focus area(s): Human-Centred AI;

Supervisors: Associate Professor Elizabeth Stephens, Dr. Bonnie Evans

Catching Prompting Bias Live: How People Use AI During Cross-Ideological Vaccination Debates

Thomas Felesina

School of Psychology, Business School

Presentation abstract: When people encounter views that challenge their own, do they seek to understand — or to win? This project introduces a novel real-time method for observing how individuals interrogate AI during ideologically charged conversations. We pair participants with opposing views on childhood vaccination for a 20-minute text-based debate via a custom-built platform. In one version of the study, each participant also has access to a private AI assistant that only they can see, allowing us to capture exactly how they prompt and use AI in the heat of a live disagreement. By logging every query directed to the AI alongside the unfolding human conversation, we can examine whether people use AI to seek disconfirming evidence, to strengthen their existing position, or to craft more persuasive arguments — effectively capturing confirmation bias as it happens. We also measure whether AI-assisted conversations produce different patterns of attitude change, moral conviction, and outgroup warmth compared to unassisted debates. The study contributes to emerging questions about how AI tools mediate human reasoning in polarised contexts — a question with growing real-world relevance as people increasingly consult AI during everyday disagreements. This presentation outlines the study design, the platform architecture, and the analytic strategy ahead of data collection.

Presentation keywords: AI-assisted debate, confirmation bias, vaccine hesitancy, attitude change, human-AI interaction

Focus area(s): Human-Centred AI;

Supervisors: Brendan Zietsch, Matthew Hornsey, Sam Pearson

Audio Retrieval: Challenges and Solutions

Watheq Mansour

EECS

Contact: https://au.linkedin.com/in/watheq-mansour

Presentation abstract: Information Retrieval systems are becoming increasingly multimodal, and one interesting and natural modality for human information interaction is audio. Audio data is rapidly expanding in both volume and popularity. For example, as of September 2025, there are 4.52 million podcasts available worldwide, with over 584 million listeners, indicating a global growth in attention to this industry. However, several challenges arise when searching this vast space of data, including content diversity, length variability, and various indexing considerations such as retrieval segmentation and text vs audio retrieval modalities. In this work, we aim to address some of these challenges and provide a useful base solution to the information retrieval community interested in this domain. We plan to achieve this goal through three main milestones. In the first one, we develop and group all relevant text-based techniques and resources. This includes developing baseline retrieval systems, and surveying and expanding existing test collections. In the second project, we conduct a systematic failure analysis of the systems resulting from the first project. The output is intended to highlight common failure cases and provide practical recommendations to avoid them. Finally, we plan to devise novel approaches benefiting from the attained recommendations, develop audio-based solutions, and explore the effect of fusing the best of both modalities, i.e., text and audio.

Presentation keywords: audio retrieval, multimodal retrieval, automatic speech recognition

Focus area(s): Data-Centric AI;

Supervisors: Joel Mackenzie, Shane Culpepper

List of Posters

AI in the detection of preclinical and early osteoarthritis

Jessica Perchman

EECS

Contact: https://www.linkedin.com/in/jess-perchman-b8bb49367/

Presentation abstract: Knee osteoarthritis (OA) affects more than a 100 million people globally and is a major cause of pain and disability costing Australia 2 billion dollars annually. AI applied to medical imaging technologies (Xray and MRI) could allow clinicians to detect, quantify and monitor early cartilage changes and enable treatment and strategies for OA before structural and functional deficits become irreversible. This project focusses on the development and validation of deep learning methods for automated cartilage segmentation and biomarker extraction from knee MRI scans. By analysing structural biomarkers such as cartilage thickness and volume, alongside compositional biomarkers such as T2 relaxation times, this research could identify areas of the joint that show preclinical signs of OA. By automating these processes, the time and expertise to extract manually segment images and extract biomarkers could be reduced significantly, enabling large-scale studies across different cohorts and support earlier diagnosis and monitoring of OA progression.

Presentation keywords: Medical Imaging, Segmentation, Biomarkers, Deep Learning, Osteoarthritis

Focus area(s): Human-Centred AI;

Supervisors: Shakes Chandra, Jurgen Mejan-Fripp, Kieran O'Brien

AI Parent Selection: Sensory and agronomic traits in strawberry (Fragaria x ananassa)

Laura Esquivel Garcia

QAAFI

Presentation abstract: Breeding for flavour in the cultivated strawberry (Fragaria × ananassa) is challenged by a complex octoploid genome and the high environmental sensitivity of sensory traits. Traditional parent selection relies on resource-intensive, multi-year field trials and trained sensory panels, which can introduce selection bias. To overcome these limitations, our study explores a framework that combines AI-driven digital twins with Fastack (haplotype stacking) to optimize parental selection of desirable flavour profiles. This approach incorporates volatile aroma compound data previously captured via high-throughput LC-MS phenotyping. Digital twins function as virtual replicas of specific strawberry genotypes, simulating how complex metabolic and volatile-production pathways respond to diverse environmental conditions. By linking these volatile aroma compound models with genomic data, we aim to establish a framework for future in-silico crosses to model haplotype stacking. This approach aims for precise targeting and accumulation of multi-locus chromosomal blocks associated with superior sweetness and aroma, while simultaneously minimizing linkage drag.

Presentation keywords: AI-Driven Digital Twins, Genomic Selection, High-Throughput Phenotyping, Plant Breeding

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

Supervisors: Heather Smyth

Ethics in the Age of AI

Michael Tafe

Business School

Contact: https://www.linkedin.com/in/michaeltafe/

Presentation abstract: Our distinct evolutionary identity as 'rational animals' appears destined to be supplanted by exponentially more capable thinking machines. AI seems set to first supersede, and then make redundant, our intellectual capability and capacity. It seems to me an urgent task to understand and safeguard what it will mean to be human – to stay human – in this dawning new Age.

Charles Darwin once surmised "Of all the differences between man and the lower animals, (it is) the moral sense or conscience (that) is by far the most important…It is the most noble of all attributes of man." But today our understanding of morals and ethics remains fractured and incomplete. Philosophy has been the traditional custodian of ethical theory but mostly prefers moralising over empirical study. And contemporary sciences such as evolutionary biology, moral psychology, and moral anthropology, advance alternate sets of detailed claims which each compete with the other for descriptive influence.

Combining theoretical and empirical knowledge from across these and related fields, my research seeks to resolve the multiplicity of narratives and provide an essential synthesis. The intention is to produce a unifying account of 'how ethical views are formed' - a descriptive, operational explanation of 'how' ethical views come to exist, be maintained and changed in the world - so that potential impacts and future opportunities for the practice of ethics in the Age of AI can be pre-empted today.

Presentation keywords: Ethics, Artificial Intelligence (AI), Human identity, Moral and Ethical Practice

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

Supervisors: Prof Thomas Maak, Prof Nicola Pless, Dr Caitlin Curtis

Enhancing Lexical Resources for Argumentative Essay Writing through Corpus and AI Integration

Nguyen Thi Thuy Linh

School of Languages and Cultures

Contact: https://www.linkedin.com/in/nguyen-thi-thuy-linh-17941037b/

Presentation abstract: This proposed study explores the potential of integrating corpus consultation and AI-mediated interaction to support the development of academic collocational competence in second language (L2) writing. The study involves approximately 80 Vietnamese undergraduate students majoring in English enrolled in academic writing courses at a public university in Vietnam. A pre-intervention, while-intervention, immediate post-intervention, and delayed post-intervention instructional design will be employed. During the intervention, students will use a corpus tool (CorpusMate) alongside structured interaction with ChatGPT to identify, analyse, and refine academic collocations in their writing. Data will be collected through collocational-competence tests, academic essay tasks, post-intervention questionnaires, and group discussions. Quantitative analyses will examine changes in learners' collocational competence, as well as the frequency, accuracy, and range of academic collocations used in their writing, while qualitative data will provide insights into learners' perceptions and experiences of using corpus tools and AI-mediated support. The study is expected to contribute to current discussions on AI-assisted language learning and offer pedagogical implications for integrating corpus tools and generative AI into L2 academic writing instruction.

Presentation keywords: Academic collocational competence, corpus-based language learning, academic writing, Generative AI in language learning

Focus area(s): Human-Centred AI;

Supervisors: Associate Professor Peter Crosthwaite, Dr Martin Schweinberger

Generating High-Fidelity Semi-Synthetic Concept Drift Streams for Electronic Medical Records

Oliver Pienaar

Queensland Digital Health Centre

Presentation abstract: Background: Deployed clinical machine-learning models are susceptible to concept drift—changes in underlying data distributions over time. Drift can degrade performance, necessitating robust monitoring and adaptation strategies. Objective: Evaluating drift detectors requires accurate datasets tailored to deployment scenarios. Because privacy constraints prevent sharing longitudinal Electronic Medical Record (EMR) data, practitioners must rely on synthetic or semi-synthetic datasets. Methods: Purely synthetic datasets often lack the complexity of real-world clinical data. We developed a novel dataset generator that constructs realistic, semi-synthetic drifting streams using public EMR data (MIMIC-IV). It injects controlled artificial drift into the static dataset, preserving the within-sample realism while creating between-sample temporal changes. Results: Experiments show our tool successfully produces high-fidelity drift streams from static EMR data. It permits precise control over drift characteristics, allowing researchers to simulate specific clinical shifts. Conclusion: By combining real-world EMR data with synthetic drift mechanics, this generator provides a vital testbed to stress-test clinical drift detectors prior to real-world deployment.

Presentation keywords: concept drift, data streams, dataset generator, clinical AI

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

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

SAMRI: Segment Any MRI — A Foundation Model for Whole-Body MRI Segmentation

Zhao Wang

EECS

Contact: www.linkedin.com/in/zhao-wang-39710025a

Presentation abstract: Magnetic resonance imaging (MRI) provides exceptional soft-tissue detail without radiation, making it indispensable across neurology, oncology, and musculoskeletal imaging. Automatically identifying organs and tissues in MRI scans (segmentation) is critical for measuring disease, planning treatment, and monitoring progress — yet remains hard to automate.

Task-specific AI models such as U-Net perform well in medical image segmentation, but must be retrained for each new organ or scan type, struggle on unseen data, and lack interactivity — a key need in clinical practice. The Segment Anything Model (SAM) addresses these gaps with a prompt-based, interactive design, but existing medical adaptations ignore MRI-specific issues like uneven brightness, blurred boundaries, and tiny structures.

My research introduces SAMRI (Segment Any MRI), an AI model tailored for whole-body MRI segmentation. A two-stage training strategy — precomputing image features once, then fine-tuning only the segmentation component — reduces training time by 94% and computational cost by 99%, enabling fine-tuning on standard hardware. SAMRI was trained on 1.1 million labelled MRI slices covering 47 organ targets across 10+ scan types.

SAMRI achieved a segmentation accuracy (mean Dice score) of 0.87 across all 47 targets and maintained strong performance on six entirely unseen datasets — demonstrating that targeted, efficient training can deliver a reliable open-source tool for research and clinical use.

Presentation keywords: MRI, Segmentation, SAM, Vision Transformer, Deep learning

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

Supervisors: Shakes Chandra