


NUS AI Research Day 2026

Technical Programme

Event Details

Date	Tuesday 22 September 2026
Time	11:00 am – 4:00 pm
Venue	COM3 Multi-Purpose Hall (#01-27) and COM3 Seminar Room 14 (#01-23)
Address	COM3 Building, 11 Research Link, Singapore 119391

Please refer to the programme below for the exact location of each session.



Multi-Purpose Hall, Level 1, COM3 Building: poster and demo layout.

Programme at a Glance

Time	Activity	Venue
11:00 am – 12:15 pm	Opening Remarks	COM3 Seminar Room 14 (#01-23)
	Faculty Spotlight Session 1*	
12:15 pm – 12:45 pm	Lunch Break	COM3 Level 1 Foyer
12:45 pm – 3:00 pm	Poster Session	COM3 Multi-Purpose Hall (#01-27)
	Demo Session	COM3 Level 1 Foyer
3:00 pm – 4:00 pm	Faculty Spotlight Session 2*	COM3 Seminar Room 14 (#01-23)
	Closing Remarks	

* livestream available in COM3 Multi-Purpose Hall (#01-27)

Faculty Spotlight Session 1

Venue: COM3 Seminar Room 14 (#01-23)

Session Chair: Yisong Miao, Siqi Ge

11:00 – 11:05 am Opening Remarks ---- Prof. Lee Wee Sun	
Faculty Spotlight Session 1	
11:05 – 11:20 am	Asst. Prof. Peichen Zhong Presidential Young Professor Department of Materials Science and Engineering, College of Design and Engineering Title: AI for Science in atomistic simulations for materials and chemistry Abstract: Materials modeling with atomistic simulation offers a powerful alternative to conventional trial-and-error methods by providing atomic-level insights for mechanistic understanding and accelerated material discovery. A central objective is to predict material properties at finite temperatures, which requires a computationally efficient statistical mechanics approach that maintains ab-initio accuracy. In this talk, I will provide an overview of how machine learning and AI-assisted models can be integrated into atomistic simulations, how these approaches can be applied to understand phenomena in materials and chemistry, and how to scale up to device-level simulations to advance technological applications.
11:20 – 11:35 am	Asst. Prof. Tan Zhi Xuan Presidential Young Professor Department of Computer Science, School of Computing Title: Scaling Rational Cooperative Intelligence for Pluralistic AI Futures Abstract: The future of AI is likely to be a pluralistic one, with a plethora of AI agents, systems and services deployed by different actors towards different interests. AI could thus accelerate harmful competition and conflict, but also enable better coordination and cooperation across society. How can we shape AI development towards the latter? To address this challenge, I argue that we need to scale cooperative intelligence, advancing the cooperative capacities and dispositions of the agents we create, while creating institutions and incentives that promote AI and human cooperation at scale. To develop these capacities, I will introduce recent research

	from my lab on rational cooperative AI — a paradigm for building coherent AI agents that are cooperative-by-design, and for designing rational incentives for AI to cooperate at scale.
11:35 am – 11:50 am	Asst. Prof. Jiang Wenqi Department of Computer Science, School of Computing Title: Efficient Computing Systems for Machine Learning Abstract: Computing infrastructure underpins advances in machine learning, and improving its efficiency is increasingly important as AI workloads grow in scale and complexity. In this talk, I will first present ROSA, a robotics foundation model serving system that enables fleets of robots to share a pool of server-class GPUs for efficient and reliable operation in robot factories. ROSA supports multi-model pipelines spanning action generation, planning, safety checking, and progress monitoring, and coordinates resource allocation to maximize factory-level productivity while meeting task-specific performance requirements. Evaluations on a real robot and large-scale simulated workloads show up to 12× higher factory productivity than conventional systems that dedicate a GPU to each robot. Then, I will briefly review my broader research on retrieval-augmented generation, efficient vector search, and data systems on modern hardware. This work spans algorithms, systems, and hardware, showing how coordinated design across the computing stack can improve the efficiency of complex AI applications.
11:50 am – 12:05 pm	Asst. Prof. Doudou Zhou Department of Statistics and Data Science, Faculty of Science Title: Learning and Evaluation with Limited Supervision: AI for Healthcare and Language Model Evaluation Abstract: Health records and AI-generated outputs offer abundant information, but trusted labels and human assessments are often scarce. My research develops statistical methods that use this information while accounting for differences across sources and errors in automated signals. I will begin with work on integrating medical concepts across hospitals by combining clinical records, medical knowledge, and language models. I will then discuss how inexpensive automated scores can improve estimates of task-specific accuracy, and how historical human comparisons can support the evaluation of newly introduced models. These problems require distinguishing the quantity we want to measure from the imperfect signals available to estimate it, while accounting for uncertainty in the resulting evaluations. I will close with future directions in allocating limited human evaluation effort, deciding when to acquire additional information, and drawing statistical conclusions from data collected through adaptive decisions.
12:05 pm – 12:15 pm	Asst. Prof. David P. Daniels Presidential Young Professor Department of Management and Organisation, NUS Business School Title: How AI Can Improve Human Decision Making in Negotiations Abstract: Decades of research on human decision making have produced few reliable ways to prevent mistakes, and no widely recognized tool for improving human decision making in negotiations or other business settings. This talk will present SAIBORGS (Science-based AI-powered Behavior-Optimizing Research and Guidance System), an AI-powered negotiation platform with three components. First, a new AI-powered sparring partner ("AI NegotiationBot"): it pushes back, adapts, and tries to influence human negotiators, offering people a way to practice negotiation on demand. Second, a new AI-powered guidance system ("AI NudgeBot"): it watches a negotiation unfold, intercepts the negotiator's intended moves in real time, and

	catches mistakes before they are made, not after, when it is too late. Third, a new AI-powered science engine ("AI ScienceBot"): it generates new negotiation data using AI NegotiationBots and/or collects data from human negotiators, and then analyzes every message and every move, making it possible to automatically test which interventions improve negotiation outcomes, and why.
--	--

Lunch Break

12:15 pm – 12:45 pm	Lunch Break Venue: Level 1 Foyer, COM3 Building
---------------------	--

Poster and Demo Session

Time: 12:45 pm – 3:00 pm

Venue: COM3 Multi-Purpose Hall (#01-27)

Demonstrations

[D1] InTriage: Intelligent Telephone Triage in Pre-Hospital Emergency Care

He Kai — Saw Swee Hock School of Public Health

Pre-hospital Emergency Care (PEC) systems are critical for managing life-threatening emergencies where rapid intervention can significantly impact patient outcomes. The rising global demand for PEC services, coupled with increased emergency calls and strained emergency departments, necessitates efficient resource utilization through Telephone Triage (TT) systems. However, existing TT processes face challenges such as incomplete data collection, communication barriers, and manual errors, leading to high over-triage and under-triage rates. This study proposes InTriage, an AI-driven multilingual TT system to provide decision support for triage. InTriage enhances accuracy by transcribing emergency calls, extracting critical patient information, prompting for supplementary information, and providing real-time triage decision support. We conducted an evaluation on a real-world corpus of approximately 40 hours of telephone data, achieving a word error rate of 14.57% for speech recognition and an F1 score of 73.34% for key information extraction. By improving communication efficiency and reducing triage errors, InTriage offers a scalable solution to potentially help address the growing demands on PEC systems globally.

Theme: AI + Healthcare and Human Potential

[D2] DeepConnect: A Visual Analytics System for Bridging Interdisciplinary Research Collaborations (IEEE VIS 2026)

Feng Yingchaojie — School of Computing

An LLM-augmented visual analytics system for interdisciplinary collaborator discovery. DeepConnect translates collaboration ideas into domain-specific tasks, retrieves relevant papers to ground cross-domain exploration, and provides coordinated visualisations for comparing candidate researchers. It also surfaces terminology gaps and overlaps across domains and supports publication-grounded conversation rehearsal for outreach.

Theme: AI + Science

[D3] TSAutoBox: A Localised Multi-Agent System for Industrial Time-Series Task Automation (VLDB 2026)

Xinyu Xi, Yihao Ang, Yifan Bao, Ye Xu, Ruiyao Ma, Anthony K. H. Tung — Department of Computer Science

An interactive, edge-deployed multi-agent system that turns a natural-language task description and a dataset into an executable time-series modelling pipeline. TSAutoBox profiles the data, infers the task, constructs an initial ToolChain, and refines candidates through search-tree expansion with deterministic execution and multi-

fidelity evaluation. Demonstrated on manufacturing anomaly detection, cybersecurity forecasting, and healthcare time-series generation, running on a local edge device.

Theme: *Agentic AI, Artificial General Intelligence, and Autonomous Systems*

[D4] Detecting Misleading Half-Truths: Diverse Evidence Retrieval and Multi-Agent Reasoning

Tang Yixuan — School of Computing

An interactive demonstration for detecting misleading half-truths caused by omitted context. Visitors can browse and search claims, inspect diverse retrieved evidence, compare presented against omitted information, and follow a role-anchored multi-agent debate to a final True / Half-True / False assessment. The system integrates NEWSCOPE for diverse evidence retrieval, TRACER for consequential omission detection, and RADAR for evidence-grounded multi-agent reasoning (EMNLP 2025, EMNLP 2026).

Theme: *Computer Vision, Natural Language Processing, and Multimodality*

[D5] Mars Compute: Hardware-Aware Agentic Optimisation for LLM Training and Inference on Heterogeneous AI Accelerators

Zhuobin Huang — School of Computing

An agentic system that automatically improves the training and inference efficiency of LLM workloads on GPU and TPU accelerators. Mars Compute combines instruction-level hardware knowledge across roughly two dozen architectures with a full-stack analysis toolkit, then runs a closed loop that profiles a workload, reasons over the results, edits the kernel, and gates every change behind performance and accuracy checks. Across 20 production kernels on four GPU platforms, it reports a 1.65× geometric-mean speedup over library baselines.

Theme: *Resource-Efficient AI and AI Hardware*

[D6] AI Tutor: A Platform for AI-Supported Professional Learning (FinTutor and InsurTutor)

Zhou Yi — Asian Institute of Digital Finance

A reusable platform for AI-supported professional learning, currently piloted in a professional certification programme for working adults in Singapore's financial sector. A curriculum knowledge graph defines concepts and prerequisites; retrieval-grounded tutoring answers strictly from approved sources with guardrails for out-of-scope queries; Bayesian knowledge tracing estimates per-concept mastery; and a reinforcement learning engine sequences questions to reach mastery faster. Visitors can question the tutor, inspect retrieved sources, take an adaptive practice set, test the guardrails, and view the learning analytics dashboard.

Theme: *AI + Education, Humanities and Social Sciences*


[D7] AutoOPH: A Multi-Agent System for Multimodal Ophthalmic Research

Wu Jianzhi — Yong Loo Lin School of Medicine

Glaucoma research requires more than fitting a model to a single image set. A study may involve fundus photographs, optical coherence tomography, structured examinations, and longitudinal records, while its conclusions depend on choices about the cohort, analysis unit, data split, and evaluation. Existing medical agents mainly support diagnosis or individual clinical tasks, whereas general research agents are not designed specifically for glaucoma data. We introduce AutoOPH, a locally deployable multi-agent system that carries a glaucoma study from literature search and data assessment to experiment design, implementation, reflection, review, and reporting. Its meta-learning strategy adapts search priorities from earlier accepted research trajectories. Also, our controller can reject a candidate with higher score after the reviewer detected reuse of the evaluation set, and the SEC prevented the rejected result from entering confirmatory memory. In demonstration glaucoma subtype prediction, AutoOPH achieved a patient-level out-of-fold macro-F1 of 0.5925.

Theme: *Multi-Agent System in Healthcare, Multimodal*

Posters



Foundational AI

Agentic AI, AI, and Autonomous Systems

- F1.** FineVerify: Scaling Test-Time Compute with Fine-Grained Self-Verification for Agentic Search, EMNLP 2026, August 2026.
[James Xu Zhao](#) — School of Computing
- F2.** EUBRL: Epistemic Uncertainty Directed Bayesian Reinforcement Learning, ICLR 2026, April 2026.
[Jianfei Ma](#) — School of Computing
- F3.** MCPMark: A Benchmark for Stress-Testing Realistic and Comprehensive MCP Use, ICLR 2026, April 2026.
[Ye Yaoqi](#) — School of Computing
- F4.** Med-Banana: Learning Quality-Controlled Medical Image Editing from Success-and-Failure Trajectories, EMNLP 2026.
[Zhihui Chen](#) — Saw Swee Hock School of Public Health
- F5.** MemEvolve: Meta-Evolution of Agent Memory Systems, ICML 2026.
[Zhang Guibin](#) — School of Computing
- F6.** Event-B Agent: Towards LLM Agent for Formal Model Synthesis and Repair, FSE 2026, July 2026.
[Hongshu Wang](#) — School of Computing
- F7.** LoopNav: Benchmarking Spatial Consistency in World Models, ICML 2026 CTB workshop, July 2026.
[Lian Kewei](#) — School of Computing

AI Theory

- F8.** The Proxy Presumption: From Semantic Embeddings to Valid Social Measures, ACL 2026.
[Kelvin Koa](#) — Asian Institute of Digital Finance
- F9.** Small coresets via negative dependence: DPPs, linear statistics, and concentration, NeurIPS 2024.
[Tran Hoang Son](#) — Department of Mathematics

F10. Last-Iterate Convergence of Randomized Kaczmarz and SGD with Greedy Step Size, COLT 2026, Jun 2026.
[Dong Xiaoyu](#) — [Department of Mathematics](#)

Computer Vision, Natural Language Processing, and Multimodality

F11. Camellia: Benchmarking Cultural Biases in LLMs for Asian Languages, EMNLP 2026, Oct 2026.
[Sahajpreet Singh](#) — [School of Computing](#)

F12. Relevance as a Vulnerability: How Web Retrieval Degrades Safety Alignment in LLM Agents, EMNLP 2026, August 2026.
[Aditya Nawal](#) — [Department of Electrical and Computer Engineering](#)

F13. Mitigating Memory Interference in Lifelong Model Editing through Multi-Head Hopfield Retrieval, EMNLP 2026, October 2026.
[Xuming Ran](#) — [Yong Loo Lin School of Medicine](#)

F14. UC-VLM: Consistency-Driven Learning for AI-Generated Image Detection with Vision-Language Large Models, ECCV 2026, September 2026.
[Lei Tan](#) — [Centre for Trusted Internet and Community](#)

F15. CLaRE-ty Amid Chaos: Quantifying Representational Entanglement to Predict Ripple Effects in LLM Editing, ACL 2026, July 2026.
[Manit Baser](#) — [Department of Electrical and Computer Engineering](#)

F16. E-TIDE: Fast, Structure-Preserving Motion Forecasting from Event Sequences, IEEE/RSJ IROS 2026, September 2026.
[Biswadeep Sen](#) — [School of Computing](#)

F17. The Devil is in the Spurious Correlations: Boosting Moment Retrieval with Dynamic Learning, ICCV 2025.
[Zhou Xinyang](#) — [Department of Electrical and Computer Engineering](#)

F18. ReasonMap: Towards Fine-Grained Visual Reasoning from Transit Maps, CVPR 2026, June 2026.
[Sicheng Feng](#) — [Department of Electrical and Computer Engineering](#)

F19. Communicating Chess Strategies in Natural Language, EMNLP 2026, October 2026.
[Cui Langyuan](#) — [School of Computing](#)

F20. Listen, Pause, and Reason: Toward Perception-Grounded Hybrid Reasoning for Audio Understanding, ACL 2026, Jul 2026.
[Wang Jieyi](#) — [Institute of Data Science](#)

F21. Rapid Soil Classification using Computer Vision with CNN, 9ICEG 2023.
[Aw You Jin Eugene](#) — [Department of Civil and Environmental Engineering](#)

Machine Learning, Deep Learning, and Foundation Models

F22. Causal Foundation Models for Time Series via Prior-Data fitted Networks, ICML 2026, May 2026.
[Dennis Thumm](#) — [Asian Institute of Digital Finance](#)

F23. JAVELIN: Directional Time Series Editing via Retrieval-Guided Jacobian-Vector Inference, KDD 2026, Aug 2026.
[Yifan Bao](#), [Yihao Ang](#) — [School of Computing](#)

F24. Why Ask One When You Can Ask k? Learning-to-Defer to the Top-k Experts, ICLR 2026.
[Yannis Montreuil](#) — [School of Computing](#)

F25. FEAT-KD: Learning Concise Representations for Single and Multi-Target Regression via TabNet Knowledge Distillation, ICML 2025, July 2025.
[Fong Kei Sen](#) — [Department of Electrical and Computer Engineering](#)

F26. Tab-PET: Graph-Based Positional Encodings for Tabular Transformers, AAAI 2026, January 2026.
[Yunze Leng](#) — [Department of Electrical and Computer Engineering](#)

F27. Escaping the Homophily Trap: A Threshold-free Graph Outlier Detection Framework via Clustering-guided Edge Reweighting, ICLR 2026, April 2026.
[Jinyu Cai](#) — [Institute of Data Science](#)

- F28.** MoST: Mixing Speech and Text with Modality-Aware Mixture of Experts, ICML 2026, April 2026.
[Yuxuan Lou — School of Computing](#)
- F29.** Letting Trajectories Spread: Quality-Preserving Control for Diverse Flow Matching, ICML 2026.
[Wan Zhenglin — School of Computing](#)
- F30.** Chronoepilogi: Scalable time series selection with multiple solutions, NeurIPS 2024.
[Vassilis Christophides — Image and Pervasive Access Laboratory \(IPAL\)](#)
- F31.** Hierarchical Decision Making with Structured Policies: A Principled Design via Inverse Optimization, ICML 2026, July 2026.
[Wang Yuexuan — Institute of Operations Research and Analytics](#)
- F32.** Practical Global and Local Bounds in Gaussian Process Regression via Chaining, AAAI 2026.
[Junyi Liu — School of Computing](#)

Reasoning AI

- F33.** An Enigma of Artificial Reason: Investigating the Production-Evaluation Gap in Large Reasoning Models.
[Sun Ming Zhong — School of Computing](#)
- F34.** X-RAY: Mapping LLM Reasoning Capability via Formalized and Calibrated Probes, KDD 2026.
[Yufan Cai — School of Computing](#)
- F35.** SPARROW: Scalable Taxonomy Induction via Structure-Preserving Partitioning and Constraint-Guided Merging, EMNLP 2026, September 2026.
[Yirui Zhang — School of Computing](#)

Resource-Efficient AI and AI Hardware

- F36.** HiT: A Unified Sparsity-Adaptive Architecture for High-Throughput Matrix Multiplication, ISCA 2026.
[Xiang Tingting — School of Computing](#)
- F37.** Topology-Mapping Co-Design for Scalable LLM Accelerators with Hierarchical Mesh-of-Trees NoC, ISCAS 2026.
[Zhen Wu — Department of Electrical and Computer Engineering](#)
- F38.** PRIMAL: Processing-In-Memory Based Low-Rank Adaptation for LLM Inference Accelerator, ISCAS 2026, May 2026.
[Yue Jiet Chong — Department of Electrical and Computer Engineering](#)
- F39.** SHIELD: A Segmented Hierarchical Memory Architecture for Energy-Efficient LLM Inference on Edge NPUs, AICAS 2026, Sept 2026.
[Zhang Jintao — Department of Electrical and Computer Engineering](#)
- F40.** LOKI: An LLM Accelerator with Optimized KV Cache and In Memory Computing Cores, ISCAS 2026, May 2026.
[Avadh Harkishanka — Department of Electrical and Computer Engineering](#)
- F41.** HiVec: Scalable and Energy-Efficient CGRA in a RISC-V SoC for Wearables, HotChips 2026.
[Rohan Juneja — School of Computing](#)
- F42.** Analysis of Memory-Runtime Trade-offs in Caching Strategies for Genetic Programming Symbolic Regression, GECCO 2025, July 2025.
[Shi Jiaming — Department of Electrical and Computer Engineering](#)
- F43.** Slow-Fast Inference: Training-Free Inference Acceleration via Within-Sentence Support Stability, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2026 (under minor revision).
[Yu Zhaochen — School of Computing](#)

Responsible and Safe AI

- F44.** Fairness in Aggregation: Optimal Top-k and Improved Full Ranking, ICML 2026, July 2026.
[Yan Hong Yao Alvin — School of Computing](#)
- F45.** Test-Time Scaling in Reasoning Models Is Not Effective for Knowledge-Intensive Tasks Yet, COLM 2026, July 2026.

[Zhao Xu — School of Computing](#)

- F46.** From Backward Spreading to Forward Replay: Revisiting Target Construction in LLM Parameter Editing, ICML 2026, July 2026.

[Liu Wei — School of Computing](#)

- F47.** RBCBF: Decoding Time Safety Alignment via Risk Guided Rollback and Barrier Control, ICML 2026, April 2026.

[Chen Tianxiang — Department of Civil and Environmental Engineering](#)

- F48.** Zombie Agents: Persistent Control of Self-Evolving LLM Agents via Self-Reinforcing Injections, ICLR 2026 Workshop on Lifelong Agents, April 2026.

[Xianglin Yang — School of Computing](#)

- F49.** Image Can Bring Your Memory Back: A Novel Multi-Modal Guided Attack against Image Generation Model Unlearning, ICLR 2026.

[Renyang Liu — Institute of Data Science](#)

- F50.** The Consent Problem for Organoid-Based Biocomputing, Nature, 30 July 2026.

[Insoo Hyun — Yong Loo Lin School of Medicine, Centre for Biomedical Ethics](#)

- F51.** CASCADE Against Jailbreaks: Combination Across Stages with Controlled Attack-Defense Evaluation, EMNLP 2026, Aug 2026.

[Luo Jiale — School of Computing](#)

- F52.** AliMark: Enhancing Robustness of Sentence-Level Watermarking Against Text Paraphrasing, ICML 2026.

[Yuexin Li — School of Computing](#)

AI + X

AI + Ageing

- X53.** AI and Elderly: Co-Designing Meaningful Human–AI Collaboration in Ageing Contexts, ECAI 2025, October 2025.

[Goh Lay Xuan Estherine — College of Design and Engineering](#)

AI + Education, Humanities and Social Sciences

- X54.** Multilingual Agent-Based World Modeling for Social Science, COLM, Oct 2026.

[Xuan Zhang — Institute of Data Science](#)

- X55.** Using Large Language Models to Design Personalized Decoy Nudges, Theory and Practice in Marketing Conference, TPM 2026, June 2026.

[Vladimir Maksimenko — Department of Civil and Environmental Engineering](#)

- X56.** Unlocking the Archive, DH 2026, July 2026.

[Jose Miguel Vieira — Department of Philosophy](#)

- X57.** ScholAlstic: Enabling Research on Human–AI Interaction in Education, June 2026.

[Karina Yuen — AI Centre of Educational Technologies](#)

AI + Finance and Economics

- X58.** QO-Bench: Diagnosing Query-Operator-Preserving Retrieval over Typed Event Tuples, EMNLP 2026 Findings.

[Tan Tianhui — Asian Institute of Digital Finance](#)

- X59.** FinStressTS: A Parametric Synthetic Benchmark for Time-Series Forecasting in Finance, KDD 2026, August 2026.

[Kelvin Koa — Asian Institute of Digital Finance](#)

- X60.** FinDeepForecast: A Live Multi-Agent System for Benchmarking Deep Research Agents in Financial Forecasting, EMNLP 2026, October 2026.

[Xuan Yao — Asian Institute of Digital Finance](#)

- X61.** CrossAlpha: An Annual-Report Benchmark for Cross-Market Factor Research (with LLM Agents), EMNLP 2026.

[Qian Wang — School of Computing](#)

- X62.** CryptoTrade: A Reflective LLM-based Agent to Guide Zero-shot Cryptocurrency Trading, EMNLP 2024, Nov 2024.

[Bingqiao Luo — Asian Institute of Digital Finance](#)

AI + Healthcare and Human Potential

- X63.** Brain-Computer Interfaces at a Crossroads: A Critical Analysis of Technical, Ethical, and Societal Hurdles on the Path to a Connected Future, Human Brain, 2026 (accepted).

[Rosalie Elvira — National Neuroscience Institute](#)

- X64.** AI Agent for Delirium Screening among Patients with Cancer: A Proof-of-Concept Study, International Surgical Oncology Symposium (29 August 2026).

[Zeng Yingchun — Alice Lee Centre for Nursing Studies, Yong Loo Lin School of Medicine](#)

- X65.** CodeBrain: Bridging Decoupled Tokenizer and Multi-Scale Architecture for EEG Foundation Model, ICLR 2026.

[Wu Feng — Saw Swee Hock School of Public Health](#)

AI + Logistics, Manufacturing and Robotics

- X66.** Prior-Knowledge-Guided Transformer-Model-Based Synthesis for Active S-Parameter Extrapolation in Linear Antenna Arrays, IEEE Transactions on Antennas and Propagation, May 2026.

[Peiqin Liu — Department of Electrical and Computer Engineering](#)

- X67.** A Pseudo-Boolean Encoding Method for Rectilinear Steiner Arborescence in VLSI Routing, ISCAS 2025.

[Zixuan Wang — Department of Electrical and Computer Engineering](#)

- X68.** Insights Beneath the Surface: Physics-guided Multi-task Learning for In Situ Defect Detection in Laser Powder Bed Fusion, Journal of Manufacturing Processes, July 2026.

[Zhang Jiayi — Department of Mechanical Engineering](#)

- X69.** Metaheuristic Color Scheme Strategies for Enhanced Visual Inspection of 3D IC Layouts, PAINE 2025, October 2025.

[Shi Zhansen — College of Design and Engineering](#)

- X70.** DBT-Bleed: Dual Branch Temporal Modeling with Key-Frame Selection for Surgical Bleeding Detection, MICCAI 2026, September 2026.

[Sudhanshu Mishra — Department of Electrical & Computer Engineering](#)

AI + Science

- X71.** Transformers with Physics-Informed Encodings and Simulation-Based Inference for Robust Detection of Eccentric Binary Black Holes in Pulsar Timing Array Data, AI4X Accelerate Conference 2026, June 2026.

[Subhajit Dandapat — Department of Physics](#)

- X72.** PPLM: A protein pair language model for protein-protein interaction modeling, Nature Communications, March 2026.

[Jun Liu — Cancer Science Institute of Singapore](#)

- X73.** Optimizing Biophysical Large-Scale Brain Circuit Models With Deep Neural Networks, Nature Methods, July 2026 (accepted, in press).

[Zeng Tianchu — Department of Electrical and Computer Engineering](#)

- X74.** OmniClimate-TC: Physics-Aware Visual Abstractions for Multimedia Reasoning over Tropical Cyclones, ACM MM 2026, November 2026.

[Luwei Xiao — College of Design and Engineering](#)

X75. Molecular Atlas of Key Food Odorants Reveals Mixture-Level Organization and Enables Generative Aroma Design, Advanced Science, 2026.

[Zhang Jingzhi](#) — Department of Food Science and Technology

AI + Urban Sustainability

X76. Understanding urban sustainability trade-offs in Singapore using Graph Neural Network with explainable AI, WCS 2026, April 2026.

[Xia Ye](#) — Department of the Built Environment

X77. Inferring Individual Travel Heterogeneity from Trajectory Maps: An MLLM-Assisted Trip Destination Choice Model, TRB, Jan 2027.

[Dan Zhu](#) — Department of Civil and Environmental Engineering

X78. Adapting Bioacoustic Classification Models to Local Soundscapes via Spatiotemporal Routing, ACM SIGSPATIAL, November 2026.

[Rhett Michael Olson](#) — Department of Computer Science

X79. AI Assisted Generative Design of Graphene-Enhanced Architected Cementitious Microstructures, AI4X 2026, June 2026.

[Abhimanyu Goel](#) — Department of the Built Environment, Institute for Functional Intelligent Materials

AI Governance, Policy, and Societal Impact

X80. LLM DNA: Tracing Model Evolution via Functional Representations, ICLR 2026, April 2026.

[Zhaomin Wu](#) — School of Computing

Early-Career Faculty Spotlight Session 2 and Closing Remarks

Venue: COM3 Seminar Room 14 (#01-23)

Session Chair: Yisong Miao, Siqi Ge

Faculty Spotlight Session 2	
3:00 pm – 3:15 pm	Asst. Prof. Alexander Vinogradov Presidential Young Professor Department of Pharmacy and Pharmaceutical Sciences, Faculty of Science Title: Accurate models of enzymatic substrate preferences and their use in drug discovery and chemical biology Abstract: This talk will focus on our recently developed integrated experimental and computational platform for studying substrate preferences of catalytically promiscuous post-translational modification enzymes at ultra-high throughput. Using mRNA display as the experimental engine generates large, high-quality datasets, which are then leveraged by deep learning methods to build accurate models of enzymatic substrate fitness landscapes. These models enable the prediction and exploration of enzyme–substrate specificity and are being applied to drug discovery, with prospective applications in chemical biology.
3:15 pm – 3:30 pm	Asst. Prof. Gianmarco Mengaldo Department of Mechanical Engineering, College of Design and Engineering Title: From Physics to AI and back Abstract: In this talk, we explore how we can use pre-existing knowledge (e.g., physics) to improve AI systems, and how we can possibly extract some knowledge from AI systems. On the first topic, we present a novel physics-enhanced deep learning hybrid method, namely CondensNet, for resolving cloud physics in climate

	models. On the second topic, we present some results on the use of explainable AI for climate applications. We conclude with an overall perspective bridging the two topics.
3:30 pm – 3:45 pm	Asst. Prof. Dr Tham Yih Chung Presidential Young Professor Department of Ophthalmology, Yong Loo Lin School of Medicine Title: Transforming Eye Care with AI: From Foundation Models to Real-World Implementation Abstract: Using eye care as a case study, this talk will explore the journey of medical AI from building foundation models, to rigorous evaluation across diverse settings, and ultimately real-world implementation. It will highlight key lessons in developing robust, clinically meaningful AI and translating technological advances into sustainable improvements in healthcare delivery.
3:45 pm – 4:00 pm	Adj. Asst. Prof. Gao Yujia Consultant, Division of Hepatobiliary & Pancreatic Surgery, Department of Surgery, National University Hospital Consultant, Adult Liver Transplantation Programme, National University Centre for Organ Transplantation, National University Hospital Title: Connected Intelligence — From Smart Hospital to Intelligent Ecosystems
4:00 pm	Closing Remarks – Prof. Mohan Kankanhalli

Best Poster Presenter Award

Foundational AI winner

- F11.** Camellia: Benchmarking Cultural Biases in LLMs for Asian Languages, EMNLP 2026, Oct 2026.
[Sahajpreet Singh — School of Computing](#)
- F23.** JAVELIN: Directional Time Series Editing via Retrieval-Guided Jacobian-Vector Inference, KDD 2026, Aug 2026.
[Yifan Bao, Yihao Ang — School of Computing](#)

AI + X Winner

- X62.** CryptoTrade: A Reflective LLM-based Agent to Guide Zero-shot Cryptocurrency Trading, EMNLP 2024, Nov 2024.
[Bingqiao Luo — Asian Institute of Digital Finance](#)
- X80.** LLM DNA: Tracing Model Evolution via Functional Representations, ICLR 2026, April 2026.
[Zhaomin Wu — School of Computing](#)

Programme Organising Committee

Advisor

Prof Lee Wee Sun, Professor in Department of Computer Science (School of Computing)
Dr Yongkang Wong, Senior Research Fellow, School of Computing [Trustworthy Multimodal AI](#)

Chair

Dr Fakhar Abbas Senior Research Fellow, CTIC [GenAI, AI regulation](#)

Committee Members

Dr Harry Cheng	Research Fellow, School of Computing	Deepfakes, ML
Miao Yisong	PhD Student, School of Computing	NLP
Lu Jiahao	PhD Student, School of Computing	
Isaac Lee Ning	PhD Student, School of Computing	Model Interpretability
Fu Guoji	PhD Student, School of Computing	
Ge Siqi	PhD Student, Duke-NUS Medical School	
Feng Yu	PhD Student, School of Computing	
Siyuan Yao	PhD Student, School of Computing	