NUS AI Research Day

Event Details

Date: Friday July 25, 2025
Time: 1:00 pm – 5:00 pm
Venue: Level 1 of COM3 Building
11 Research Link, Singapore 119391
Please see program below for exact locations for the various sessions.

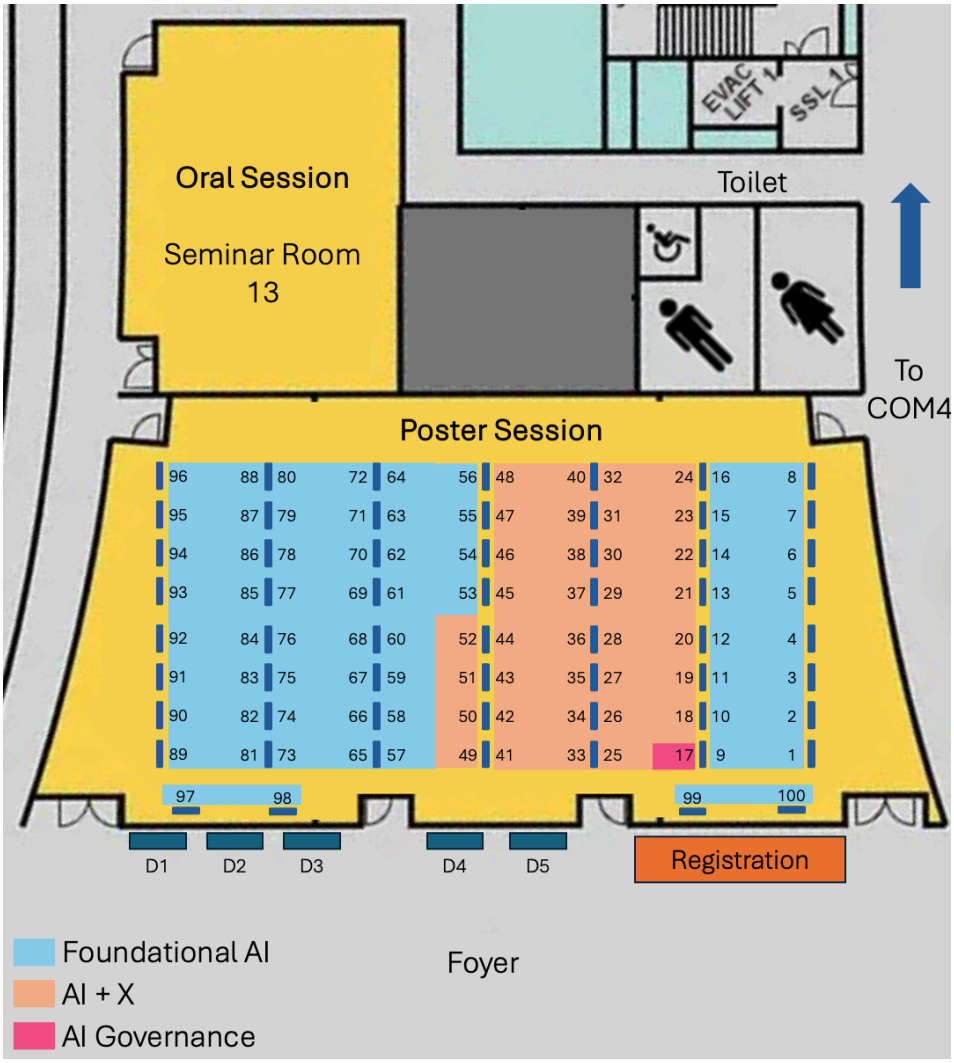

Oral Session

Oral Session 1 Venue: COM3 Seminar Room 13 (#01-22) Session Chair: Zhang Chi	
1:30 pm – 1:45 pm	From Weight-Based to State-Based Fine-Tuning: Further Memory Reduction on LoRA with Parallel Control, ICML 2025 *Chi Zhang, REN Lianhai, Jingpu Cheng, Qianxiao Li* Resource-Efficient AI
1:45 pm – 2:00 pm	Stochastic Taylor Derivative Estimator: Efficient amortization for arbitrary differential operators, NeurIPS 2024 *Zekun Shi, Zheyuan Hu, Min Lin, Kenji Kawaguchi* AI + Science
2:00 pm – 2:15 pm	AlphaEdit: Null-Space Constrained Knowledge Editing for Language Models, ICLR 2025 (Outstanding Paper) Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, Tat-Seng Chua AI Governance

Oral Session 2 Venue: COM3 Seminar Room 13 (#01-22) Session Chair: Ziwei Xu	
3:30 pm – 3:45 pm	Time will Tell: Lightweight Time Series Analytics with Low Data Requirement *Yihao Ang, Qiang Huang,Yifan Bao, Anthony K. H. Tung, Zhiyong Huang* Generative AI
3:45 pm – 4:00 pm	LivePoem: Improving the Learning Experience of Classical Chinese Poetry with AI-Generated Musical Storyboards, IJCAI 2025 *Qihao Liang, Xichu Ma, Torin Hopkins, Ye Wang* AI + Education, Humanities and Social Sciences
4:00 pm – 4:45 pm	Building Trust in the Generative AI Era: A Systematic Review of Global Regulatory Frameworks to Combat the Risks of Mis-, Dis-, and Mal-Information *Fakhar Abbas* AI Governance

Poster and Demo Sessions

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

Demo

[D1] **A Plug-and-Play AI Box for Just-in-Time Decision Intelligence**

Just-in-Time.AI offers a plug-and-play AI box for just-in-time decision intelligence, addressing the critical need for early anomaly detection, a common yet challenging requirement across manufacturing, finance, cybersecurity, and healthcare, complicated further by scarce data on rare or special cases. Combining patented anomaly detection with controllable and interpretable data generation, it empowers organizations to detect risks early, simulate edge scenarios, clearly explain model decisions, and proactively respond in real-time, all within a lightweight, edge-ready deployment.

[D2] **CogniVision**

CogniVision is the first zero-wire AI vision system-on-chip with >10-year battery life. It is the outcome of an NRF-sponsored research project, and uniquely covers the entire signal chain from sensing to AI. Its extremely low consumption enables the sense of vision anywhere (ubiquitous), all the time (always-on).

[D3] **NewsKaleido**

NewsKaleido is a Chrome extension designed to complement your current news reading by helping you discover relevant yet diverse articles. It leverages diversified retrieval and interpretable representation learning to provide a more comprehensive understanding of the news. Explore multiple perspectives, reduce bias, and make more informed decisions.

[D4] **Streamlining Contract Life Cycle Management with Advanced AI**

We exploit the power of GenAI to manage contracts more easily and smartly at scale. In our demo, we will show that we can chat with contracts, auto fill reports, review the contracts, and highlight any legal risks, etc.

[D5] **ScholAIstic**

ScholAIstic is an innovative way for teaching human skills through generative AI. Instructors carefully design realistic scenarios for learners to navigate complex situations in a safe environment. ScholAIstic allows instructors to curate conversations that promote pedagogical integrity. Some use cases include role-playing personas for competency-based learning, learning assistants that cultivate self-directed learning, and agentic facilitators that guide students through advanced cognitive tasks such as research and analysis.

Poster

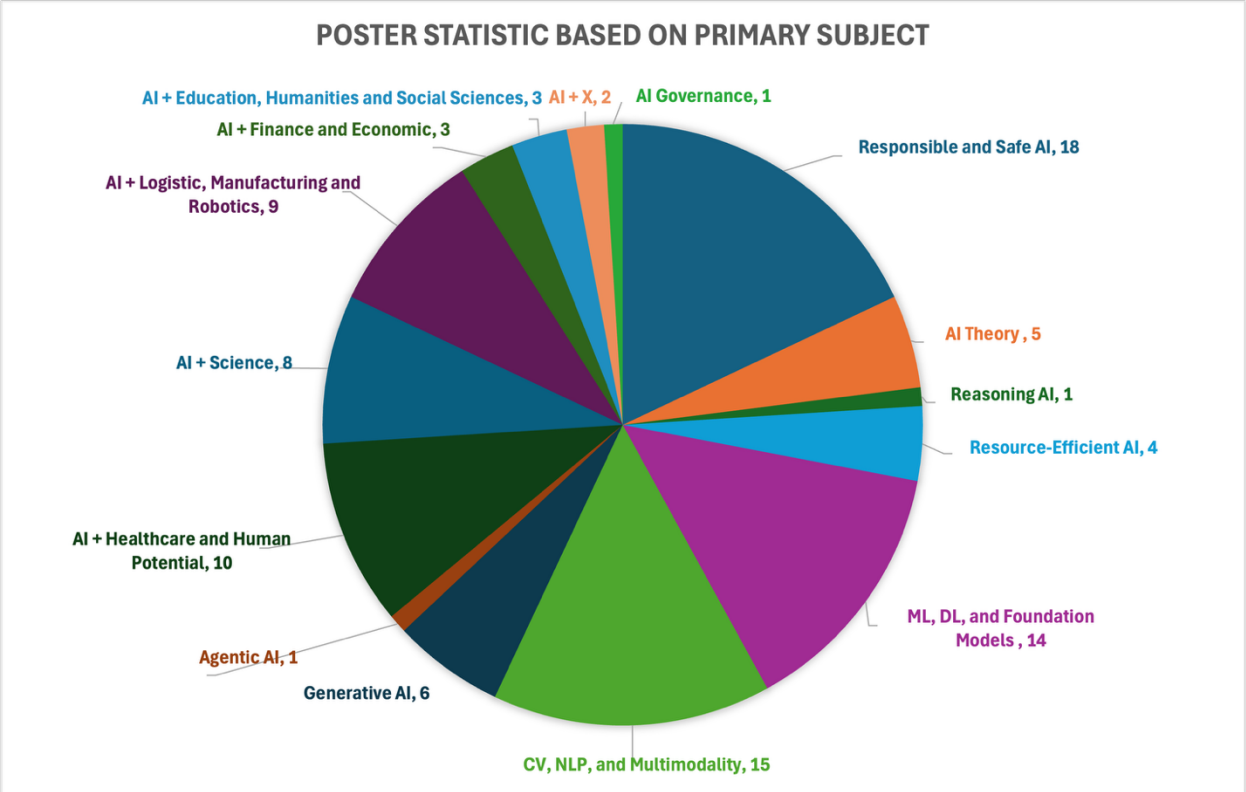

Agentic AI

1. ShowUI: One Vision-Language-Action Model for GUI Visual Agent, CVPR 2005.
   Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Weixian Lei, Lijuan Wang, Mike Zheng Shou

Computer Vision, Natural Language Processing, and Multimodality

2. CAMPrints: Leveraging the "Fingerprints" of Digital Cameras to Combat Image Theft, MobiSys 2025.
   Bangjie Sun, Mun Choon Chan, Jun Han

3. NightHaze: Nighttime Image Dehazing via Self-Prior Learning, AAAI 2025.
   Beibei Lin, Yeying Jin, Wending Yan, Wei Ye, Yuan Yuan, Robby T. Tan

4. Teaching AI the Anatomy Behind the Scan: Addressing Anatomical Flaws in Medical Image Segmentation with Learnable Prior, ICCV 2025.
   Young Seok Jeon, Hongfei Yang, Huazhu Fu, Mengling Feng

5. EventFly: Event Camera Perception from Ground to the Sky, CVPR 2025.
   Lingdong Kong, Dongyue Lu, Xiang Xu, Lai Xing Ng, Wei Tsang Ooi, Benoit R. Cottereau

6. F^3Set: Towards Analyzing Fast, Frequent, and Fine-grained Events from Videos, ICLR 2025.
   Zhaoyu Liu, Kan Jiang, Murong Ma, Zhe Hou, Yun Lin, Jin Song Dong

7. WASA: Watermark-based Source Attribution For Large Language Model-Generated Data, ACL 2025.
   Xinyang Lu, Jingtan Wang, Zitong Zhao, Zhongxiang Dai, Chuan-Sheng Foo, See-Kiong Ng, Bryan Kian Hsiang Low
8. GLOBAL Encoding and Controlling Global Semantics for Long-form Video Question Answering, EMNLP 2024.
   Thong Nguyen, Zhiyuan Hu, Xiaobao Wu, Cong-Duy T Nguyen, See-Kiong Ng, Luu Anh Tuan
9. Visual Intention Grounding for Egocentric Assistants, ICCV 2025
   Pengzhan Sun, Junbin Xiao, Tze Ho Elden Tse, Yicong Li, Arjun Akula, Angela Yao
10. Efficient Zero-Shot HOI Detection: Enhancing VLM Adaptation with Innovative Prompt Learning, NeurIPS 2024.
    Qinqian Lei, Bo Wang, Robby T. Tan
11. HOLa: Zero-Shot HOI Detection with Low-Rank Decomposed VLM Feature Adaptation, ICCV 2025.
    Qinqian Lei, Bo Wang, Robby T. Tan
12. Learning 4D Panoptic Scene Graph Generation from Rich 2D Visual Scene, CVPR 2025.
    Shengqiong Wu, Hao Fei, Jingkang Yang, Xiangtai Li, Juncheng Li, Hanwang Zhang, Tat-seng Chua
13. Synthetic-to-Real Self-supervised Robust Depth Estimation via Learning with Motion and Structure Priors, CVPR 2025.
    Weilong Yan, Ming Li, Haipeng Li, Shuwei Shao, Robby T. Tan
14. Domain-Adaptive 2D Human Pose Estimation via Dual Teachers in Extremely Low-Light Conditions, ECCV 2024.
    Yihao Ai, Yifei Qi, Bo Wang, Yu Cheng, Xinchao Wang, Robby T. Tan
15. FractalForensics: Proactive Deepfake Detection and Localization via Fractal Watermarks, ACM Multimedia 2025.
    Tianyi Wang, Harry Cheng, Ming-Hui Liu, Mohan Kankanhalli
16. Mamba as a Bridge: Where Vision Foundation Models Meets Vision Language Models for Domain-Generalized Semantic Segmentation, CVPR 2025.
    Xin Zhang, Robby T. Tan

AI Governance

17. Unmasking deepfakes: A systematic review of deepfake detection and generation techniques using artificial intelligence, Expert Systems with Applications 2024.
    Fakhar Abbas, Araz Taeihagh, Simon Chesterman

AI + Healthcare and Human Potential

18. Deep learning modelling to forecast emergency department visits using calendar, meteorological, internet search data and stock market price, Computer Methods and Programs in Biomedicine July 2025.

Chua Ming, Geraldine JW Lee, Yao Neng Teo, Yao Hao Teo, Xinyan Zhou, Elizabeth SY Ho, Emma MS Toh, Marcus Eng Hock Ong, Benjamin YQ Tan, Andrew FW Ho

19. Insights Into the Current and Future State of AI Adoption Within Health Systems in Southeast Asia: Cross-Sectional Qualitative Study, Journal of Medical Internet Research, June 2025.
Mochammad Fadjar Wibowo, Alexandra Pyle, Emma Lim, Joshua W Ohde, Nan Liu, Jonas Karlström

20. ST-USleepNet: A Spatial-Temporal Coupling Prominence Network for Multi-Channel Sleep Staging, IJCAI 2025.
Jingying Ma, Qika Lin, Ziyu Jia, Mengling Feng

21. Identifying who are unlikely to benefit from total knee arthroplasty using machine learning models, NPJ Digital Medicine.
Xiaodi Liu, Yingnan Liu, Mong Li Lee, Wynne Hsu, Ming Han Lincoln Liow

22. Scene-to-Audio: Distant Scene Sonification for Blind and Low Vision People, CHI 2025.
Chitralekha Gupta, Ashwin Ram, Shreyas Sridhar, Christophe Jouffrais, Suranga Nanayakkara

23. VRSense: An Explainable System to Help Mitigate Cybersickness in VR Games, CHI 2025.
Dinithi Dissanayake, Chitralekha Gupta, Prasanth Sasikumar, Suranga Nanayakkara

24. Exploiting Negative Samples: A Catalyst for Cohort Discovery in Healthcare Analytics, ICML 2024.
Kaiping Zheng, Horng-Ruey Chua, Melanie Herschel, H. V. Jagadish, Beng Chin Ooi, James Wei Luen Yip

25. Smart Imitator: Learning from Imperfect Clinical Decisions, Journal of the American Medical Informatics Association.
Dilruk Perera, Siqi Liu, Kay Choong See, Mengling Feng

26. WalkingWizard - A Truly Wearable EEG Headset for Everyday Use, ACM Transactions on Computing for Healthcare 2023.
Teck Lun Goh, Li-Shiuan Peh

27. ProMind-LLM: Proactive Mental Health Care via Causal Reasoning with Sensor Data, ACL 2025 Finding.
Xinzhe Zheng, Sijie Ji, Jiawei Sun, Renqi Chen, Wei Gao, Mani Srivastava

AI + Science

28. Physics-informed machine learning for MPC: Application to a batch crystallization process.
Guoquan Wu, Zhe Wu

29. Two-Stage Optimization for Ammonia Combustion Mechanisms with Improved Deep Reinforcement Learning, ASPACC 2025.
Hao Hu, Huangwei Zhang

30. Acoustic insulation optimization of walls and panels with functional graded hollow sections using graph transformer evaluator and probability-informed genetic algorithm, Building and Environment, Vol 270, 15 February 2025, 112550.

Hanmo Wang, Tam H Nguyen, Zhuyin Lu, Febi Rhiana, Dickson Fong, Wei He, Shawn Owyong, Huan Ting Chen, Alexander Lin

31. Biosensing Using CNNs to Detect Noisy But Persistent Nanoparticle Binding Events on Supported LIPID Bilayers, AI4X 2025.
Zhen Yuan Yeo, Eun Ho Song, Kyeongmee Lee, Jwa-Min Nam, N. Duane Loh

32. AI-driven High-quality structure modeling for protein monomer and complex.
Wei Zheng, Qiqige Wuyun, Yang Li, Lydia Freddolino, Yang Zhang

33. US-align: Accurate and fast structure alignments of macromolecular structures.
Chengxin Zhang, Morgan Shine, Lydia Freddolino, Anna Marie Pyle, Yang Zhang

34. NExT-Mol: 3D Diffusion Meets 1D Language Modeling for 3D Molecule Generation, ICLR 2025.
Zhiyuan Liu, Yanchen Luo, Han Huang, Enzhi Zhang, Sihang Li, Junfeng Fang, Yaorui Shi, Xiang Wang, Kenji Kawaguchi, Tat-Seng Chua

35. Language-Enhanced Representation Learning for Single-Cell Transcriptomics, Arxiv 2025.
Yaorui Shi, Jiaqi Yang, Changhao Nai, Sihang Li, Junfeng Fang, Xiang Wang, Zhiyuan Liu, Yang Zhang

AI + Logistic, Manufacturing and Robotics

36. Prior-Knowledge-Guided Deep-Learning-Enabled (PK-DL) Methods for Antenna Optimization and Synthesis.
Zhi Ning Chen, Peiqin Liu, Zijue Shan, Ziqi Zhu

37. A Graph Convolutional LSTM Approach for Modeling Nonlinear Chemical Process Networks Using Spatial-Temporal Data, Comp. & Chem. Eng.,109242, 2025.
Wanlu Wu, Guoquan Wu, Zhe Wu

38. FLIP: Flow-Centric Generative Planning For General-Purpose Manipulation Tasks, ICLR 2025.
Chongkai Gao, Haozhuo Zhang, Zhixuan Xu, Cai Zhehao, Lin Shao

39. From Detection to Forecasting: Leveraging Time-Series Foundation Models for Defects Anticipation in Metal Additive Manufacturing, Journal of Manufacturing Processes.
Jiayi Zhang, Farzam Farbiz, Mehdi Jafary-Zadeh, Swee Leong Sing

40. Chemotactic navigation in robotic swimmers via reset-free hierarchical reinforcement, Nature Communications.
Tongzhao Xiong, Zhaorong Liu, Yufei Wang, Chong Jin Ong, Lailai Zhu

41. Revealing the underlying mechanism in controlling Young's modulus of additively manufactured Ti-6Al-4V using fuzzified machine learning, Virtual and Physical Prototyping.
Yanting Liu, Cherq Chua, Verner Soh, Zhongji Sun, Chee Kai Chua, Swee Leong Sing

42. MARL-enhanced Federated Learning for Predicting Remaining Useful Lifetime, VTC 2025 Spring.
Rajarshi Chattopadhyay and Chen-Khong Tham

43. Biomimetic self-shading walls via 3D-printing for reduced heat gain: Multiscale learning using graph neural networks to predict solar radiation absorption, Building and Environment, Vol 279, 1 July 2025, 113048.
Hanmo Wang, Zhuyin Lu, Amanda Wojtasiak, Shawn Owyong, Bo Sun, Yu Fang, Shahin Vassigh, Alexander Lin

44. $D(R, O)$ Grasp: A Unified Representation of Robot and Object Interaction for Cross-Embodiment Dexterous Grasping, ICRA 2025.
Zhenyu Wei, Zhixuan Xu, Jingxiang Guo, Yiwen Hou, Chongkai Gao, Zhehao Cai, Jiayu Luo, Lin Shao

AI + Finance and Economics

45. Risk-Sensitive Deep Reinforcement Learning for Portfolio Optimization, Journal of Risk and Financial Management.
Xinyao Wang, Lili Liu

46. Algorithm Aversion: Evidence from Ridesharing Drivers, Management Science.
Meng Liu, Xiaocheng Tang, Siyuan Xia, Shuo Zhang, Yuting Zhu, Qiannying Meng

47. Through the Looking Glass: Leverage Machine Learning to Price Corporate Carbon Footprints, CICF 2025.
Haoxu Wang, Ahmed Syalabi Seet, Zhou Zhou, Johan Selaeman

AI + Education, Humanities and Social Sciences

48. Probing Historical Image Contexts: Enhancing Visual Archive Retrieval through Computer Vision, ACM Journal on Computing and Cultural Heritage.
Lin Du, Brandon Le, Edouardo Honig

49. GenAI for Creativity and Ethical Use: Enhancing Creativity and Higher-Order Thinking in Design Thinking Courses.
Kate Sangwon Lee, Li Hong Idris Lim

50. CustomAIzEd: Bridging Interdisciplinary Gaps in AI Education with Customised Content Using LLMs, IEEE EDUCOM 2025.
Li Xue Pua, Rushil Ramesh, Prabhu Natarajan, Ganesh Neelakanta Iyer

AI + X

51. AI-Enhanced Physical Layer for Next-Gen Wireless Networks
Kanav Sabharwal, Mun Choon Chan

52. TE-TM Balanced Wide-Angle Metacells for Low Scan-Loss Metalens Antenna Using Prior Knowledge-Guided Generative Deep Learning-Enabled Method, IEEE Transactions on Antennas and Propagation, May 2025.
Yanhe Lyu, Zhi Ning Chen

Machine Learning, Deep Learning and Foundation Models

53. Omni-Mol: Exploring Universal Convergent Space for Omni-Molecular Tasks, ICLR 2025 LMRL Workshop.
Chengxin Hu, Hao Li, Yihe Yuan, Zezheng Song, Haixin Wang

54. A dual-aspect evaluation framework for architectural-like plan generation via pix2pix series algorithms, Frontiers of Architectural Research.
Yu Guo, Tianyu Fang, Zhe Cui, Rudi Stouffs

55. Symbolic Regression Enhanced Decision Trees for Classification Tasks, AAAI 2024.
Kai Sen Fong, Mehul Motani

56. Language Representations Can be What Recommenders Need: Findings and Potentials, ICLR 2025.
Leheng Sheng, An Zhang, Yi Zhang, Yuxin Chen, Xiang Wang, Tat-Seng Chua

57. Understanding the Role of Initialization Schemes for State Space Models, ICLR 2025.
Fusheng Liu, Qianxiao Li

58. Cross-Domain Feature Augmentation for Domain Generalization, IJCAI 2024.
Yingnan Liu, Yingtian Zou, Rui Qiao, Fusheng Liu, Mong Li Lee,Wynne Hsu

59. Reward Shaping for Reinforcement Learning with An Assistant Reward Agent, ICML 2024.
Haozhe Ma, Kuankuan Sima, Thanh Vinh Vo, Di Fu, Tze-Yun Leong

60. DUPRE: Data Utility Prediction for Efficient Data Valuation, AAMAS 2025.
Kie Thao Nguyen Nguyen, Rachael Hwee Ling Sim, Quoc Phong Nguyen, See-Kiong Ng, Bryan Kian Hsiang Low

61. NICE Data Selection for Instruction Tuning in LLMs with Non-differentiable Evaluation Metric, ICML 2025.
Jingtan Wang, Xiaoqiang Lin, Rui Qiao, Pang Wei Koh, Chuan-Sheng Foo, Bryan Kian Hsiang Low

62. Unsupervised Machine Learning for Phase Identification and Characterization of High-Resolution STEM EELS in Novel Battery Materials, APMC 2025.
Zhen Yuan Yeo, Wenhui Lai, Jong Hak Lee, Deepan Balakrishnan, Barbaros Ozyilmaz, N. Duane Loh

63. TOPA: Extending Large Language Models for Video Understanding via Text-Only Pre-Alignment, NeurIPS 2024.
Wei Li, Hehe Fan, Yongkang Wong, Mohan Kankanhalli, Yi Yang

64. Mutual Information-Based Evolutionary Feature Construction via Minimizing Redundancy and Maximizing Relevance, GECCO 2025.
Yunze Leng, Kei Sen Fong, Mehul Motani

65. SkyLadder: Better and Faster Pretraining via Context Window Scheduling
Tongyao Zhu, Qian Liu, Haonan Wang, Shiqi Chen, Xiangming Gu, Tianyu Pang, Min-Yen Kan

66. Broaden your SCOPE! Efficient Multi-turn Conversation Planning for LLMs with Semantic Space, ICLR 2025.
Zhiliang Chen, Xinyuan Niu, Chuan-Sheng Foo, Bryan Kian-Hsiang Low

Reasoning AI

67. An Adaptive Kernel Approach to Federated Learning of Heterogeneous Causal Effects, NeurIPS 2022.
Thanh Vinh Vo, Arnab Bhattacharyya, Young Lee, Tze-Yun Leong

Resource-Efficient AI

68. Analysis of Memory-Runtime Trade-offs in Caching Strategies for Genetic Programming Symbolic Regression, GECCO 2025.
Jiamin Shi, Kai Sen Fong, Mehul Motani

69. A Generalized Accelerator for Variably Sparse Matrix Computations in ML
Pranav Dangi, Zhenyu Bai, Rohan Juneja, Dhananjaya Wijerathne, Tulika Mitra

70. PACE: A Scalable and Energy Efficient CGRA in a RISC-V SoC for Edge Computing Applications
Vishnu P. Nambir, Yi Sheng Chong, Thilini Kaushalya Bandara, Dhananjaya Wijerathne, Zhaoying Li, Rohan Juneja, Li-Shiuan Peh, Tulika Mitra, Anh Tuan Do

71. TETRIS: Optimal Draft Token Selection for Batch Speculative Decoding, ACL 2025.
Zhaoxuan Wu, Zijian Zhou, Arun Verma, Alok Prakash, Daniela Rus, Bryan Kian Hsiang Low

Responsible and Safe AI

72. Position: Trustworthy AI Agents Require the Integration of Large Language Models and Formal Methods, ICML 2025.
Yedi Zhang, Yufan Cai, Xinyue Zuo, Xiaokun Luan, Kailong Wang, Zhe Hou, Yifan Zhang, Zhiyuan Wei, Meng Sun, Jun Sun, Jing Sun, Jin Song Dong

73. Single Character Perturbations Break LLM Alignment, AAAI 2025
Leon Lin, Hannah Brown, Kenji Kawaguchi, Michael Shieh

74. Fair Deepfake Detectors Can Generalize, Arxiv 2025.
Harry Cheng, Ming-Hui Liu, Yangyang Guo, Tianyi Wang, Liqiang Nie, Mohan Kankanhalli

75. Synthesis and Verification of String Stable Neural Certificates for Interconnected Systems, ICLR 2025; VerifiAI workshop.
Jingyuan Zhou, Haoze Wu, Longhao Yan, Kaidi Yang

76. Chat-IRB? How Application-Specific Language Models Can Enhance Research Ethics Review, BMJ Journal of Medical Ethics.
Sebastian Porsdam Mann, Joel Jiehao Seah, Stephen R. Latham, Julian Savulescu, Mateo Aboy, Brian D. Earp

77. Seeing Through Deepfakes: A Human-Inspired Framework for Multi-Face Detection, ICCV 2025.
Juan Hu, Shaojing Fan, Terence Sim

78. Towards Fair Representation: Clustering and Consensus, COLT 2025.
Diptarka Chakraborty, Kushagra Chatterjee, Debarati Das, Tien Long Nguyen, Romina Nobahari

79. Can I Hear Your Face? Pervasive Attack on Voice Authentication Systems with a Single Face Image, USENIX Security Symposium 2024.
Nan Jiang, Bangjie Sun, Terence Sim, Jun Han

80. Semantic-based Uncertainty Quantification in LLMS: A Case Study on Medical Explanation Generation Task, ICLR 2025; QUESTION Workshop.
Nicholas Kian Boon Tan, Mehul Motani

81. Waterfall: Scalable Framework for Robust Text Watermarking and Provenance for LLMs, EMNLP 2024.
Gregory Kang Ruey Lau, Xinyuan Niu, Hieu Dao, Jiangwei Chen, Chuan-Sheng Foo, Bryan Kian Hsiang Low

82. Pointwise Information Measures as Confidence Estimators in Deep Neural Networks: A Comparative Study, ICML 2025.
Shelvia Wongso, Rohan Ghosh, Mehul Motan

83. Black-Box AI and Patient Autonomy, Minds and Machines 2025.
Sinead Prince, James Edgar Lim

84. Attitudes towards the use of Artificial Intelligence in Healthcare: A Conjoint Analysis Survey in Singapore, Asian Bioethics Review.
Hui Jin Toh, Hui Yun Chan, Thomas Ploug, Søren Holm, Tamra Lysaght

85. Improved Rank Aggregation under Fairness Constraint, IJCAI 2025.
Diptarka Chakraborty, Himika Das, Sanjana Dey, Alvin Hong Yao Yan

86. GuardReasoner: Towards Reasoning-based LLM Safeguards, ICLR 2025; FM-Wild Workshop.
Yue Liu, Hongcheng Gao, Shengfang Zhai, Jun Xia, Tianyi Wu, Zhiwei Xue, Yulin Chen, Kenji Kawaguchi, Jiaheng Zhang, Bryan Hooi

87. GuardReasoner-VL: Safeguarding VLMs via Reinforced Reasoning, ICML 2025; R2-FM Workshop.
Yue Liu, Shengfang Zhai, Mingzhe Du, Yulin Chen, Tri Cao, Hongcheng Gao, Cheng Wang, Xinfeng Li, Kun Wang, Junfeng Fang, Jiaheng Zhang, Bryan Hooi

88. AlphaEdit: Null-Space Constrained Knowledge Editing for Language Models, ICLR 2025 (Outstanding Paper)
Junfeng Fang, Houcheng Jiang, Kun Wang, Yunshan Ma, Jie Shi, Xiang Wang, Xiangnan He, Tat-Seng Chua

89. Joint Vision-Language Social Bias Removal for CLIP, CVPR 2025.
Haoyu Zhang, Yangyang Guo, Mohan Kankanhalli

Generative AI

90. Expandable SPAD-based real-time gun muzzle flash localization system using FPGAs and deep-learning, Pattern Recognition and Tracking XXXIV, vol. 12527, SPIE 2023.
Generating synthetic data and training muzzle flash detection systems using GANs, Automatic Target Recognition XXXIII, vol. 12521, SPIE 2023.
Lih-Wei Chia, Mehul Motani

91. Show-o: One Single Transformer to Unify Multimodal Understanding and Generation, ICLR 2025.
Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, Mike Zheng Shou

92. CRAB: A Novel Configurable Role-Playing LLM with Assessing Benchmark, ACL 2025.
Kai He, Yucheng Huang, Wenqing Wang, Delong Ran, Dongming Sheng, Junxuan Huang, Qika Lin, Jiaxing Xu, Wenqiang Liu, Mengling Feng

93. Mercury: A Code Efficiency Benchmark for Code Large Language Models, NeurIPS 2024.
Mingzhe Du, Luu Anh Tuan, Bin Ji, Qian Liu, See-Kiong Ng

94. Fast But Accurate: A Real-Time Hyperelastic Simulator with Robust Frictional Contact, SIGGRAPH 2025.
Ziqiu Zeng, Siyuan Luo, Fan Shi

95. How Does Response Length Affect Long-Form Factuality?, ACL 2025 Finding.
James Xu Zhao, Jimmy Z.J. Liu, Bryan Hooi, See-Kiong Ng

AI Theory

96. Neural Dueling Bandits: Preference-Based Optimization with Human Feedback, ICLR 2025.
Arun Verma, Zhongxiang Dai, Xiaoqiang Lin, Patrick Jaillet, Bryan Kian Hsiang Low

97. Leveraging (Biased) Information: Multi-armed Bandits with Offline Data, ICML 2024.
Wang Chi Cheung, Lixing Lyu

98. Near Optimal Non-asymptotic Sample Complexity of 1-Identification, ICML 2025.
Zitian Li, Wang Chi Cheung

99. Quantum Physics Stochastic Neural Networks, ICSRS 2024
Umberto Alibrandi, Claudio Perez, Khalid M. Mosalam

100. Ordered V-information Growth: A New Perspective on Shared Information
Rohan Ghosh, Mehul Motani