



The Sixth International Conference on
Distributed Artificial Intelligence

Conference Manual

December 18-22, 2024
Singapore

Sponsors:



School of
Computing and
Information Systems



Introduction

International Conference on Distributed Artificial Intelligence (DAI) aims to bring together researchers and practitioners in related areas (e.g., general AI, multi-agent systems, distributed learning, computational game theory) to provide a high-profile, internationally renowned forum for research in the theory and practice of distributed AI. The conference series was initiated in 2019 in Beijing, China. Subsequently, DAI conferences have been successfully held in Nanjing, China (2020), Shanghai, China (2021) and Tianjin, China (2022), Singapore (2023). This year, the sixth edition of DAI conference will be held at Singapore Management University and have a pure format of offline in-person participation from December 18 to 22, 2024.

DAI 2024 will invite Sergey Levine, associate professor at UC Berkeley, Georgios Piliouras, research scientist at Google DeepMind, Richard S. Sutton, Professor at University of Alberta, and founder at Openmind Research Institute, Shuicheng Yan, Honorary Advisor, Kunlun Tech and Skywork AI to give keynote speeches. Beside keynote speeches, main conference, workshops and tutorials, DAI 2024 also invites papers from sister conferences (e.g., AAMAS, AAI, ICML, ICLR, NeurIPS, CoRL) accepted in past two years to present at the conference. The organization committee aims to make DAI 2024 a premier and influential venue to promote research in distributed AI and foster scientific exchange between researchers, practitioners, scientists, students, and engineers across the entirety of distributed AI and its affiliated disciplines.

Program Schedule

Time in UTC+8

Venue: SMU Administration Building

Level 4, Admin Executive Seminar Room 4-1 (AESR 4-1), Admin Executive Seminar Room 4-2 (AESR 4-2)

Level 5, Admin Mochtar Riady Auditorium (AMRA), Admin Executive Seminar Room 5-1 (AESR 5-1), Admin Executive Media Theatre (AEMT)

Dec. 18 - Day 1 - Tutorials and Workshops

Date / Time	Topic	
9:00-10:15	Tutorial 1 Diffusion Models for Decision-Making Agents (AESR 4-1)	Workshop 1 Embodied AI (AESR 4-2)
10:15-10:30	Coffee Break	
10:30-12:00	Workshop 2 Game Theory (AESR 4-1)	Workshop 1 Embodied AI (AESR 4-2)
12:00-13:30	Lunch Break	
13:30-14:45	Tutorial 2 Causal Empowered Agents & Foundation Models (AESR 4-1)	Workshop 3 Open Environment Multi-agent Reinforcement Learning (AESR 4-2)
14:45-15:00	Coffee Break	
15:00-18:00	Workshop 4 LLM-based Code Agent (AESR 4-1)	Workshop 3 Open Environment Multi-agent Reinforcement Learning (AESR 4-2)

Dec. 19 - Day 2 - AI Agent Day

Date / Time	Topic
9:00-9:45	AI Agent Day Regular Session (AMRA)
9:45-10:15	Coffee Break
10:15-12:15	AI Agent Day Regular Session (AMRA)
12:15-13:30	Lunch Break
13:30-13:40	LLM Reasoning Forum Opening (AMRA)
13:40-15:00	LLM Reasoning Forum (AMRA)
15:00-15:20	Coffee Break
15:20-16:40	LLM Reasoning Forum (AMRA)
16:40-17:30	Panel Discussion (AMRA)

Dec. 20 - Day 3 - Main Conference

Date / Time	Topic	
9:00-9:20	Opening (AMRA)	
9:20-10:20	Keynote 1 - Decentralized Neural Networks Richard S. Sutton Chair: Bo An (AMRA)	
10:20-10:50	Coffee Break	
10:50-12:30	Offline Learning and Trajectory Optimization Chair: Lei Yuan (AEMT)	Deep Learning and Application Chair: Lei Feng (AESR 5-1)
12:30-13:30	Lunch Break	
13:30-14:30	Keynote 2 - Foundations of Foundation Models Shuicheng Yan Chair: Yaodong Yang (AMRA)	
14:30-15:00	Coffee Break	
15:00-16:20	MARL and Applications I Chair: Tianpei Yang (AEMT)	MARL and Application II Chair: Youzhi Zhang (AESR 5-1)
18:30-22:30	Banquet	

Dec. 21 - Day 4 - Main Conference

Date / Time	Topic	
9:00-10:00	Keynote 3 - Robotic Foundation Models Sergey Levine Chair: Weinan Zhang (AMRA)	
10:00-10:30	Coffee Break	
10:30-11:30	Game Theory and Mechanism Design I Chair: Weiran Shen (AEMT)	Game Theory and Mechanism Design II Chair: Wanyuan Wang (AESR 5-1)
11:30-13:00	Lunch Break	
13:00-14:00	Keynote 4 - Multi-Agent Learning in Games Georgios Piliouras Chair: Yang Yu (AMRA)	
14:00-15:30	Single Agent RL and Application Chair: Chao Yu (AEMT)	

DAI-24 Conference Committee

General Co-Chairs



Sarit Kraus

Bar-Ilan University
Israel



Bo An

Nanyang Technological University
Singapore

Program Co-Chairs



Long Tran-Thanh

University of Warwick
United Kingdom



Yaodong Yang

Peking University
China

Workshop/Tutorial Co-Chairs



Bryan Wilder

Carnegie Mellon University
United States



Weinan Zhang

Shanghai Jiaotong University
China

DAI-24 Conference Committee

Sponsorship Chair



Yang Yu
Nanjing University
China

Publicity Chair



Dengji Zhao
ShanghaiTech University
China

Publication Chair




Siheng Chen
Shanghai Jiao Tong University
China

Local Chair



Xinrun Wang
Singapore Management University
Singapore

Keynote Speakers



Richard S. Sutton

Professor, University of Alberta, and
founder, Openmind Research Institute

09:20 - 10:20,
Dec. 20, 2024, UTC+8

About the Speaker

Rich Sutton, FRS, FRSC, is an artificial intelligence researcher specializing in reinforcement learning. He works at Keen Technologies, the University of Alberta, the Alberta Machine Intelligence Institute, and the Openmind Research Institute. Earlier he studied at Stanford University and the University of Massachusetts, and worked at GTE Labs, AT&T Labs, and DeepMind. His scientific publications have been cited about 150,000 times. He is also a libertarian, a chess player, and a cancer survivor.

Decentralized Neural Networks

Abstract

Artificial neural networks and deep learning, the leading edge of modern AI, are modeled after a distributed system—the biological brain—but de-emphasize its distributed aspects. For example, modern deep learning emphasizes the synchronized convergence of whole networks and gradient descent on a single global error surface. The perspective of the individual neuron-like unit is sadly almost entirely absent from deep learning's theory and algorithms. In this talk I investigate taking the decentralized perspective of the single unit to see if it can ease some of recurring problems with deep learning, such as loss of plasticity, catastrophic forgetting, and slow online learning. I highlight the recent work led by Dohare on loss of plasticity in Nature and by Elsayed on streaming reinforcement-learning algorithms.

Keynote Speakers



Shuicheng Yan

Honorary Advisor,
Kunlun Tech & Skywork AI,
Singapore
13:30 - 14:30,
Dec. 20, 2024, UTC+8

About the Speaker

Prof. Yan Shuicheng is currently the Honorary Advisor of Kunlun Tech and was formerly the Chief Scientist at Sea. He is a member of the Singapore Academy of Engineering and a Fellow of AAI, ACM, IEEE, and IAPR, and his research interests include computer vision, machine learning, and multimedia analytics. To date, Prof. Yan has published more than 800 papers in top international journals and conferences with an H-index of 140+. Prof. Yan Shuicheng's team has won more than ten awards in two core competitions, Pascal VOC and ImageNet (ILSVRC). In addition, his team has won more than ten best paper and best student paper awards, especially a grand slam at ACM Multimedia, the top conference in multimedia, including three best paper awards, two best student paper awards, and a best demo award.

Foundations of Foundation Models

Abstract

In this presentation, I divide the research on foundation models into three aspects, corresponding to three types of errors: approximation error, estimation error, and optimization error. For approximation error, I will share two recent advancements in model architecture, MoE++ and MoH. For optimization error, I will present our new optimizer called Adan and discuss how we optimize individual layers based on samples with varying levels of difficulty. Additionally, I will introduce two of our new products, SkyMusic and SkyReels.

Keynote Speakers



Sergey Levine

Associate Professor,
UC Berkeley
9:00 - 10:00,
Dec. 21, 2024, UTC+8

About the Speaker

Sergey Levine received a BS and MS in Computer Science from Stanford University in 2009, and a Ph.D. in Computer Science from Stanford University in 2014. He joined the faculty of the Department of Electrical Engineering and Computer Sciences at UC Berkeley in fall 2016. His work focuses on machine learning for decision making and control, with an emphasis on deep learning and reinforcement learning algorithms. Applications of his work include autonomous robots and vehicles, as well as applications in other decision-making domains. His research includes developing algorithms for end-to-end training of deep neural network policies that combine perception and control, scalable algorithms for inverse reinforcement learning, deep reinforcement learning algorithms, and more.

Robotic Foundation Models

Abstract

General-purpose pretrained models have transformed natural language processing, computer vision, and other fields. In principle, such approaches should be ideal in robotics: since gathering large amounts of data for any given robotic platform and application is likely to be difficult, general pretrained models that provide broad capabilities present an ideal recipe to enable robotic learning at scale for real-world applications. From the perspective of general AI research, such approaches also offer a promising and intriguing approach to some of the grandest AI challenges: if large-scale training on embodied experience can provide diverse physical capabilities, this would shed light not only on the practical questions around designing broadly capable robots, but the foundations of situated problem-solving, physical understanding, and decision making. However, realizing this potential requires handling a number of challenging obstacles. What data shall we use to train robotic foundation models? What will be the training objective? How should alignment or post-training be done? In this talk, I will discuss how we can approach some of these challenges.

Keynote Speakers



Georgios Piliouras

Research Scientist,
Google DeepMind
13:00 - 14:00,
Dec. 21, 2024, UTC+8

About the Speaker

Georgios Piliouras is game theory team lead at Google DeepMind and associate professor at the Singapore University of Technology and Design (SUTD). His research interests lie in the areas of multi-agent learning, algorithmic game theory, blockchain, and dynamical systems. He received his PhD in Computer Science from Cornell University. He has held academic positions at Georgia Institute of Technology, California Institute of Technology, and UC Berkeley. He is a research affiliate of the Ethereum Foundation. He is the recipient of a Singapore NRF Fellowship and a Simons Fellowship. His work has been recognized with multiple top honors and awards at premier conferences on Machine Learning, AI, multi-agent systems, and blockchain, including in venues such as ICLR, AAI, AAMAS, and IEEE Blockchain.

Multi-Agent Learning in Games

Abstract

We examine some classic questions in game theory and online learning. How do standard learning dynamics such as multiplicative weights update, gradient descent, and variants thereof behave when applied in games? The traditional approach to this question is to connect it to different notions of game theoretic equilibria. We discuss why such results can be uninformative in practice and instead offer new insights by using tools and techniques from dynamical systems such as chaos theory, recurrence, and conservative systems.

Workshops

Workshop 1: Embodied AI

9:00-12:00, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-2.

Organizers: Junting Chen, Siang Chen, Zhengbang Zhu, Minghuan Liu, Lin Shao

Speakers: Fan Shi, Malika Meghjani, Ziwei Wang, Hanbo Zhang, Chongkai Gao

Website: <https://sites.google.com/view/dai24-embodiedai/>

As embodied AI evolves, enabling agents to seamlessly operate in real-world, dynamic environments requires addressing complex challenges in perception, adaptation, and interaction. Unlike traditional AI models, embodied agents engage with physical spaces, necessitating resilient sensory processing and context-aware decision-making. This workshop seeks to advance these capabilities, focusing on perception-action loops and adaptive strategies that empower AI systems to handle unpredictable environments. Our aim is to bring together experts from robotics, multi-agent systems, and cognitive science to discuss foundational research and applications in areas like autonomous systems and collaborative robotics. By fostering cross-disciplinary insights, the workshop will spotlight innovations that accelerate embodied AI's practical impact.

Workshop 2: Game Theory

10:30-12:00, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-1.

Organizer: Dengji Zhao

Speakers: Chongjun Wang, Weiran Shen, Wanyuan Wang, Zhiyu Zhao

Website: <https://smart.sist.shanghaitech.edu.cn/dai24/index.html>

This workshop aims to explore the intersection of game theory and AI, focusing on how game-theoretic concepts can be applied to the design, analysis, and behavior of AI agents. We will delve into topics such as multi-agent systems, reinforcement learning, and mechanism design, with a special emphasis on large models like deep neural networks and generative models.

Workshops

Workshop 3: Multi-Agent Reinforcement Learning

13:30-18:00, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-2.

Organizers: Lei Yuan, Jianhong Wang

Speakers: Wanyuan Wang, Robert Loftin, Youzhi Zhang, David Mguni, Xingxing Liang, Dawei Qiu, Weiyu Ma, Feng Chen

Website: <https://sites.google.com/view/dai-2024-marl>

This workshop delves into the dynamic and complex world of multi-agent systems (MAS) operating in challenging environments, where solutions may involve multi-agent reinforcement learning (MARL) or other advanced techniques. Participants are invited to share case studies and real-world applications, showcasing how MAS can drive innovation in areas like embodied agents, autonomous vehicles, robotics, and more.

Workshop 4: LLM-based Code Agent

15:00-18:00, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-1.

Organizers: Yun Lin, Weinan Zhang

Speakers: Tianwei Zhang, Yi Li, Xiaofei Xie, Binhang Qi, Chenyan Liu

Website: <https://wlca2024.github.io/>

The emergence of language model changes the way of how we program. In the long run, programmers will write, edit, test, debug, and repair code with code agents. This workshops covers the cutting-edge topic on how AI and software engineering research is developed to generate, edit, test, and repair the code. By walking through all techniques to automate programmers' work, we aim to forsee how auto-programming techniques can advance our life in the upcoming 5-10 years.

Tutorials

Tutorial1 - From Theory to Practice: Diffusion Models in Sequential Decision-Making

9:00-10:15, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-1.

Abstract

Diffusion models have emerged as a powerful alternative to traditional generative models, offering improved sample quality and training stability. This tutorial provides an in-depth exploration of the application of diffusion models in enhancing sequential decision-making solutions, aiming to inspire new research in this rapidly evolving area. We begin by examining the core challenges that traditional sequential decision-making algorithms face, followed by a structured taxonomy of current methods that leverage diffusion models to address these challenges. Additionally, we highlight a series of successful real-world applications where diffusion models have been effectively integrated, showcasing their practical advantages across diverse domains.




Ting Long

Ting Long is an Associate Professor at Jilin University, specializing in reinforcement learning and information retrieval. With extensive research experience in generative models and reinforcement learning, she will cover the theoretical foundations of diffusion models and their taxonomy within decision-making. Ting Long will also discuss how these models address key challenges in decision-making contexts. Website: <https://sai.jlu.edu.cn/info/1094/4084.htm>



Zhengbang Zhu

Zhengbang Zhu is a fourth-year PhD student at Shanghai Jiao Tong University, with research expertise in imitation learning, reinforcement learning, and robotics applications. He will introduce representative algorithms in each category, providing practical insights, recommendations, and hands-on experiences for applying diffusion models in various sequential decision-making scenarios. Website: <https://zbzhu99.github.io/>



Weinan Zhang

Weinan Zhang is now a Professor (with tenure) at the department of computer science and engineering, Shanghai Jiao Tong University. His research interests include reinforcement learning, agents, and large models for decision-making. with various real-world applications of robotic control, game AI, recommender systems, etc. He has published over 200 research papers at prestigious international conferences and journals, accumulating over 20k citations on Google Scholar, been selected as Elsevier China Highly Cited Researcher. He has been serving as an area chair at ICML, NeurIPS, ICLR, KDD, etc. and an associate editor at TPAMI and FCS. Weinan served as a PC co-chair in DAI 2023. Website: <https://wnzhang.net>

Tutorials

Tutorial2 - Causality and AI Agent

13:30-14:45, Dec. 18, 2024, UTC+8. Seminar Room: Level 4, AESR 4-1.

Abstract

In AI, agent decision making has evolved from control and planning methods to advanced reinforcement learning, now including large language models (LLMs) that leverage expert data, prompt engineering, and human feedback. However, traditional approaches often fall short in dynamic, open-world settings. This tutorial advocates for causality as a critical enhancement for agent decision making, enabling a deeper understanding of environment structures and more reliable interventions. By integrating causal reasoning—through disentanglement, counterfactuals, and targeted policy analysis—LLMs can make adaptive, trustworthy decisions. We conclude by exploring challenges and future directions to further integrate causality into decision-making frameworks for robust AI systems.



Mengyue Yang

Mengyue Yang is currently the Lecturer (equals to US assistant professor) in AI at University of Bristol. She studied for her Ph.D. in Computer Science at University College London, under the supervision of Professor Jun Wang. Her main research interests include causal representation learning, multi-agent systems, and reinforcement learning, with a primary focus on decision-making systems based on causal representations. She has published several research papers in top-tier AI conferences and journals such as NeurIPS, CVPR, KDD, SIGIR, WWW, and ACM TOIS. She has been recognized as a Rising Star in AI by KAUST. Additionally, she serves as a PC member or reviewer for conferences and journals including NeurIPS, ICML, ICLR, KDD, AISTATS and TNNLS etc.

Website: <https://ymy4323460.github.io/>



Haoxuan Li

Haoxuan Li is a fourth-year Ph.D. candidate at Peking University, advised by Prof. Xiao-Hua Zhou, coadvised by Prof. Zhi Geng and Prof. Peng Cui. He has more than 30 publications appeared in several top conferences such as ICML, NeurIPS, ICLR, SIGKDD, WWW, SIGIR, AAI, and IJCAI. His research interests span from causal machine learning theory, counterfactual fairness, recommender system debiasing, out-of-distribution generalization, multi-source data fusion, bioinformatics, and large language models. Moreover, he is supported by the Young Scientists Fund of the National Natural Science Foundation of China (¥300,000), and have served as the AC or SPC/PC-member for top-tier conferences including ICML, NeurIPS, ICLR, SIGKDD, WWW, AAI, IJCAI, and the invited reviewer for prestigious journals such as TOIS, TPAMI, TKDE, TKDD, TNNLS, JASA, SCIENCE CHINA Information Sciences, and The Innovation.

AI Agent Day - Regular Session

Keynote

Building Knowledgeable Agents by Reinforcement Learning from Language Models

9:00-9:45, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.



Yang Yu

Yang Yu is a Professor at the School of Artificial Intelligence, Nanjing University. His research focuses on artificial intelligence, machine learning, and reinforcement learning. He has published over 100 papers in top-tier journals and conferences. His work has received significant recognition, including the Best Paper Award at the 2024 International Conference on Distributed Artificial Intelligence and four other international paper awards. He has also won three international algorithm competition championships, including the OpenAI Transfer Reinforcement Learning Competition.

Abstract

Large language models (LLMs) have demonstrated remarkable capacity for encoding and reasoning over vast amounts of knowledge, but their potential for decision-making remains limited due to their lack of inherent agency. While many existing approaches attempt to patch LLMs to create knowledgeable agents, an alternative path could be reinforcement learning (RL) centered. In this talk, I will explore this paradigm and present our efforts to integrate LLM knowledge into RL systems, particularly offline RL. By leveraging the knowledge-rich representations of LLMs, we enable RL agents to tackle tasks beyond their original interactive training scope while maintaining the adaptive capabilities of traditional RL systems, such as learning from feedback data. This approach opens up new opportunities for building knowledgeable agents that combine the knowledge power of LLMs with the decision-making flexibility of RL, offering a potential pathway to scalable and continually improving intelligent systems.

AI Agent Day - Regular Session

Regular Talks

10:15-12:15, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.

Hammer: Robust Function-Calling for On-Device Language Models via Function Masking

Qiqiang Lin, Muning Wen, Qiuying Peng, Guanyu Nie, Junwei Liao, Jun Wang, Xiaoyun Mo, Jiamu Zhou, Cheng Cheng, Yin Zhao, Jun Wang, Weinan Zhang

AgentBoard: An Analytical Evaluation Board of Multi-Turn LLM Agents (NeurIPS'24)

Chang Ma, Junlei Zhang, Zhihao Zhu, Cheng Yang, Yujiu Yang, Yaohui Jin, Zhenzhong Lan, Lingpeng Kong, Junxian He

TWOSOME: An Efficient Online Framework to Align LLMs with Embodied Environments via Reinforcement Learning (ICLR'24)

Weihao Tan, Wentao Zhang, Shanqi Liu, Longtao Zheng, Xinrun Wang, Bo An

Reinforcing LLM Agents via Policy Optimization with Action Decomposition (NeurIPS'24)

Muning Wen, Ziyu Wan, Weinan Zhang, Jun Wang, Ying Wen

Vision-Language-Action Models for Robot Manipulation

Yifan Zhong

TRAD: Enhancing LLM Agents with Step-Wise Thought Retrieval and Aligned Decision (SIGIR'24)

Ruiwen Zhou, Yingxuan Yang, Muning Wen, Ying Wen, Wenhao Wang, Chunling Xi, Guoqiang Xu, Yong Yu, Weinan Zhang

AI Agent Day - LLM Reasoning Forum

Advancing and Evaluating Agentic Reasoning of LLMs

13:40-14:20, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.



Junxian He

Junxian He is an assistant professor in the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology. He received his PhD degree from Carnegie Mellon University, Language Technologies Institute. He serves as the area chair of ICLR, ACL, and EMNLP. His recent research focuses on complex reasoning/planning, mechanistic interpretation, and multimodal understanding of LLMs.

Abstract

Complex reasoning is one of the most critical abilities for LLM agents, to reason, plan, and make predictions. In this talk, I will cover our recent research on advancing reasoning abilities of LLM agents for various reasoning and planning tasks. First, I will introduce B-STaR, a self-improving algorithm that balances exploration and exploitation to achieve scalable improvements for self-taught reasoners. Next, I will discuss our research on non-myopic generation, which enhances the performance of language models across multiple agentic scenarios at inference time. Lastly, I will talk about our evaluation framework, AgentBoard, to evaluate LLMs as agents in a fine-grained manner.

AI Agent Day - LLM Reasoning Forum

Agent K: Towards High-Performing Autonomous Data Science Agents

14:20-15:00, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.



Alex Maraval

Alex Maraval is a Senior ML Engineer at Huawei Noah's Ark Lab in London. He graduated from EPFL (Lausanne, Switzerland) in Mathematics and completed a Master's degree at Imperial College London in Machine Learning. He joined the Decision Making and Reasoning Team in the London Research Centre in 2020 where he started working on Variational Inference, Reinforcement Learning with a focus Gaussian Processes and Bayesian Optimization (BO). Alex contributed to a multitude of projects including research on High-Dimensional BO on structured spaces, BO on Graphs, ... He contributed to several publications in top-tier conferences including state-of-the-art algorithm HEBO and is the first author of Meta-Learning for BO with Transformer Neural Processes, published at NeurIPS 2023. More recently, Alex has been focusing on LLMs related projects. His research directions include building specialized Agents, extending RAG techniques, researching more performant optimizers and improving fine-tuning.


Abstract

Data science has long been essential, driving ongoing efforts to create agents capable of tackling complex tasks autonomously. While many such agents exist, they are often limited in scope, needing end-to-end automation or falling short in performance. In this talk, we'll explore Agent K, a new, fully autonomous agent that achieves both complete end-to-end automation and high-level performance. We created a benchmark based on real Kaggle competitions to evaluate its capabilities. Agent K achieved a 92.5% automation success rate in testing, verified through unit tests, and demonstrated its ability to handle multimodal tasks across diverse domains. Agent K consistently ranked in the top 38% among human competitors in these same competitions. Additionally, Agent K earned two bronze medals in featured competitions, making it an (unofficial) Kaggle expert. Furthermore, results comparable to achieving six gold, three silver, and seven bronze medals across all competition types—an unofficial skill set similar to Grandmaster-level medal performance.

AI Agent Day - LLM Reasoning Forum

Unlocking RL Potential: The Power of Generative Models in Complex Environments

15:20-16:00, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.



Zhongwen Xu

Dr. Zhongwen Xu is a Principal Scientist at Tencent, where his research focuses on deep reinforcement learning and generative models. He earned his Ph.D. from the University of Technology Sydney under the supervision of Prof. Yi Yang. He received his Bachelor's degree from Zhejiang University in 2013, where he was advised by Prof. Yueting Zhuang and Prof. Fei Wu. Before joining Tencent, Dr. Xu served as a Principal Scientist at Sea AI Lab, an Adjunct Assistant Professor at the National University of Singapore, and a Senior Research Scientist at DeepMind.

Abstract

This talk explores how generative models, including large language models (LLMs) and video generation models, are transforming reinforcement learning (RL) in complex 3D game environments. We focus on two key advancements: First, we demonstrate how learned world simulator models, derived from environment dynamics, significantly reduce the data needed to train effective RL agents and facilitate generalization on unseen environments. Second, we introduce a novel LLM-driven approach that streamlines the development process, enabling the creation of agents capable of playing thousands of 3D games with unprecedented efficiency.

AI Agent Day - LLM Reasoning Forum

Improve Multi-step Reasoning for LLMs with Deliberate Planning

16:00-16:40, Dec. 19, 2024, UTC+8. Seminar Room: Level 5, AMRA.



Chaojie Wang

Dr. Chaojie Wang is currently working as Research Scientist in Skywork AI 2050, which focuses on accelerating the pace of realization of Artificial General Intelligence (AGI). Before that, Chaojie obtained Ph.d degree from Xidian University in 2021 and worked as Research Fellow in Nanyang Technological University until 2023. Chaojie has published more than 30 papers in top AI conference and journals, such as T-PAMI, NeurIPS, ICML etc, leading a research team of nearly 10 members focusing on generative models, large language models (LLM), and reinforcement learning from human feedback (RLHF) technologies. In 2021, as the first contributor, Chaojie won the championship of the L2RPN-2021 international competition, and is currently focused on the development and implementation of the Skywork's supper app.

Abstract

Large Language Models (LLMs) trained on vast corpora of text data have demonstrated an impressive capability across various natural language tasks. However, the auto-regressive generation process makes LLMs prone to produce errors, hallucinations and inconsistent statements when performing multi-step reasoning, especially for solving mathematical problems and code generation. In this talk, we will focus on sharing our practical experiences in enhancing the multi-step reasoning capabilities of Skywork-o1 through cutting-edge alignment techniques, specifically Reinforcement Learning from Human Feedback (RLHF) and tree-search-based planning methods. Additionally, we will also discuss the opportunities and challenges that may be encountered in the future.

Presentation Schedule

Session: Offline Learning and Trajectory Optimization

10:50-12:20, Dec. 20, 2024, UTC+8. Seminar Room: Level 5, AEMT.

Chair: Lei Yuan

Is Mamba Compatible with Trajectory Optimization in Offline Reinforcement Learning (NeurIPS'24)

Yang Dai, Oubo Ma, Longfei Zhang, Xingxing Liang, Shengchao Hu, Mengzhu Wang, Shouling Ji, Jincai Huang, Li Shen

InPTR: Integration Prioritized Trajectory Replay (DAI'24)

Chendie Yao, Xingxing Liang, Longfei Zhang, Jincai Huang, Jun Lei, Yulong Zhang

Multi-Agent Trajectory Prediction with Scalable Diffusion Transformer (DAI'24)

Shenyu Zhang, Shixiong Kai, Chang Chen, Yuzheng Zhuang, Zhengbang Zhu, Minghuan Liu, Weinan Zhang

SCaR: Refining Skill Chaining for Long-Horizon Robotic Manipulation via Dual Regularization (NeurIPS'24)

Zixuan Chen, Ze Ji, Jing Huo, Yang Gao

Towards Smaller and Faster GPTs (DAI'24 Student Poster)

Sathya Krishnan Suresh, Shunmugapriya P

SOPPU: Scalable One PEFT per User (DAI'24 Student Poster)

Yash Jain, Mohor Banerjee

Presentation Schedule

Session: Deep Learning and Application

10:50-12:30, Dec. 20, 2024, UTC+8. Seminar Room: Level 5, AESR 5-1.

Chair: Lei Feng

Factor Graph Neural Network Meets Max-Sum: A Real-Time Route Planning Algorithm for Massive-Scale Trips (AAMAS'24)

Yixuan Li, Wanyuan Wang, Weiyi Xu, Yanchen Deng, Weiwei Wu

Moderate Message Passing Improves Calibration: A Universal Way to Mitigate Confidence Bias in Graph Neural Networks (AAAI'24)

Min Wang, Hao Yang, Jincai Huang, Qing Cheng

State Abstraction via Deep Supervised Hash Learning (TNNLS'24)

Guang Yang, Zheng Xu, Jing Huo, Shangdong Yang, Tianyu Ding, Xingguo Chen

DAG-based Column Generation for Adversarial Team Games (ICML'24)

Youzhi Zhang, Bo An, Daniel Dajun Zeng

Visual Whole-Body Control for Legged Loco-Manipulation (CoRL'24)

Minghuan Liu, Zixuan Chen, Xuxin Cheng, Yandong Ji, Ri-Zhao Qiu, Ruihan Yang, Xiaolong Wang

Presentation Schedule

Session: MARL and Applications I

15:00-16:00, Dec. 20, 2024, UTC+8. Seminar Room: Level 5, AEMT.

Chair: Tianpei Yang

Variational Stochastic Games (DAI'24)

Zhiyu Zhao, Haifeng Zhang

Boosting Studies of Multi-Agent Reinforcement Learning on Google Research Football Environment: the Past, Present, and Future (AAMAS'24)

Yan Song, He Jiang, Haifeng Zhang, Zheng Tian, Weinan Zhang, Jun Wang

TaxAI: A Dynamic Economic Simulator and Benchmark for Multi-Agent Reinforcement Learning (AAMAS'24)

Qirui Mi, Siyu Xia, Yan Song, Haifeng Zhang, Shenghao Zhu, Jun Wang

Session: MARL and Applications II

15:00-16:20, Dec. 20, 2024, UTC+8. Seminar Room: Level 5, AESR 5-1.

Chair: Youzhi Zhang

Multi-agent Multi-game Entity Transformer: Towards Generalist Models in MARL (DAI'24)

Rundong Wang, Weixuan Wang, Xianhan Zeng, Liang Wang, Zhengjie Lian, Yiming Gao, Feiyu Liu, Siqin Li, Xianliang Wang, Qiang Fu, Wei Yang, Lanxiao Huang, Longtao Zheng, Zinovi Rabinovich, Bo An

Agent-Specific Effects: A Causal Effect Propagation Analysis in Multi-Agent MDPs (ICML'24)

Stelios Triantafyllou, Aleksa Sukovic, Debmalya Mandal, Goran Radanovic

A Comprehensive Evaluation Framework for Multi-Agent Reinforcement Learning (DAI'24)

Zonglei Jing, Xiaojun Chang, Mingxuan Zhu, Simin Li, Aishan Liu, Xiaoqian Li, Xianglong Liu

SHAQ: Incorporating Shapley Value Theory into Multi-Agent Q-Learning (NeurIPS'22)

Jianhong Wang, Yuan Zhang, Yunjie Gu, Tae-Kyun Kim

Presentation Schedule

Session: Game Theory and Mechanism Design I

10:30-11:30, Dec. 21, 2024, UTC+8. Seminar Room: Level 5, AEMT.

Chair: Weiran Shen

A Game Theory Reward Model for Federated Learning with Probabilistic Verification (DAI'24)

Gennaro Auricchio, Harry Clough, Christopher Ho, Kaigui Bian, Changyu Dong, Kan Yang, Jie Zhang

Optimal Fixed-Price Mechanism with Signaling (DAI'24)

Zhikang Fan, Weiran Shen

RegFTRL: Efficient Equilibrium Learning in Two-Player Zero-Sum Games (DAI'24)

Zijian Fang, Zongkai Liu, Chao Yu

Session: Game Theory and Mechanism Design II

10:30-11:30, Dec. 21, 2024, UTC+8. Seminar Room: Level 5, AESR 5-1.

Chair: Wanyuan Wang

Efficient and Scalable Reinforcement Learning for Large-scale Network Control (NMI'24)

Chengdong Ma, Aming Li, Yali Du, Hao Dong, Yaodong Yang

Opponent Modeling in Multiplayer Imperfect-Information Games (DAI'24)

Sam Ganzfried, Kevin Wang, Max Chiswick

Safe and Robust Subgame Exploitation in Imperfect Information Games (ICML'24)

Zhenxing Ge, Zheng Xu, Tianyu Ding, Linjian Meng, Bo An, Wenbin Li, Yang Gao

Presentation Schedule

Session: Single Agent RL and Application

14:00-15:30, Dec. 21, 2024, UTC+8. Seminar Room: Level 5, AEMT.

Chair: Chao Yu

Adaptive Command: Real-Time Policy Adjustment via Language Models in StarCraft II (DAI'24)

Weiyu Ma, Dongyu Xu, Shu Lin, Haifeng Zhang, Jun Wang

Logarithmic Function Matters Policy Gradient Deep Reinforcement Learning (DAI'24)

Qi Liu, Jingxiang Guo, Zhongjian Qiao, Pengbin Chen, Jinxuan Zhu, Yanjie Li

CoMoU: Model-based Dynamics Estimation for Efficient Offline-to-online Reinforcement Learning (DAI'24)

Dongxiang Chen, Ying Wen

Looking Ahead to Avoid Being Late: Solving Hard-Constrained Traveling Salesman Problem (DAI'24)

Jingxiao Chen, Ziqin Gong, Minghuan Liu, Jun Wang, Yong Yu, Weinan Zhang

Sponsorship

Platinum Sponsors



Gold Sponsor



Special Supporter

Venue Sponsor and Administrative Support



School of
**Computing and
Information Systems**

Conference Venues

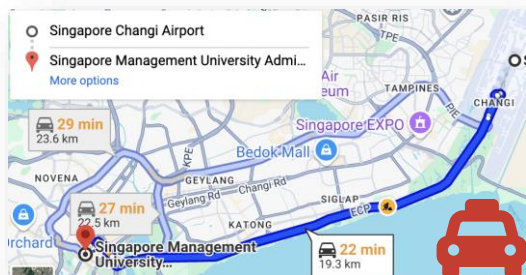
Singapore Management University

The main conference, tutorials, and workshops will be held at the Singapore Management University (SMU) 81 Victoria Street, Singapore 188065



Transport from Changi Airport to SMU

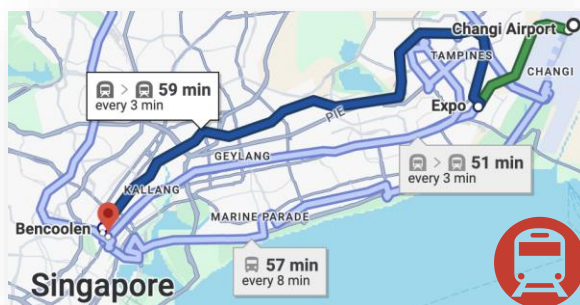
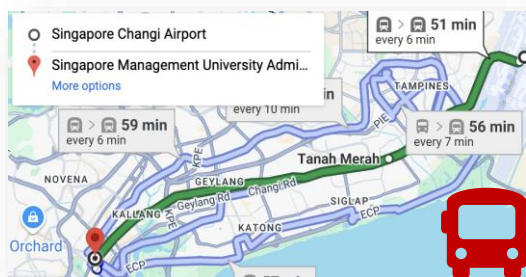
Route: Taxi from Changi Airport to SMU Admin Building.



Route: 1) Take (EW/CG) line from Changi Airport. 2) Transfer to (DT) line at Expo/(CG1/DT35). 3) Alight at Bencoolen (DT21). 4) Walk underground to Bras Basah (CC2), Exit A. 5) Walk to SMU.

Note: Nearest MRTs to SMU are Bencoolen (DT21), Bras Basah (CC2), and City Hall (EW13/NS25), all within walking distance.

Route: 1) Take Bus No.36 from T4. 2) Get off at Capitol Bldg station. 3) Walk 5 minutes to SMU.



Details: <http://www.adai.ai/dai/2024/venue.html>

Conference Rooms

SMU Administration Building

Level4

Admin Executive Seminar Room 4-1 (AESR 4-1)

Admin Executive Seminar Room 4-2 (AESR 4-2)

Level5

Admin Mochtar Riady Auditorium (AMRA)

Admin Executive Seminar Room 5-1 (AESR 5-1)

Admin Executive Media Theatre (AEMT)



Banquet

Carlton Hotel Singapore

The banquet will be held at Carlton Hotel Singapore.

76 Bras Basah Rd, Singapore 189558

Reception: 18:30-19:00, Dec. 20, 2024, UTC+8.

Banquet: 19:00-22:30, Dec. 20, 2024, UTC+8.



Transport from SMU to Banquet

It takes only a 5-minute walk to reach the banquet venue.

