Xinrun Wang

School of Computing and Information Systems
Singapore Management University

Email: xrwang@smu.edu.sg
Website: https://rainwangphy.github.io/

Research Interests

(Multi-Agent) (Reinforcement) Learning, and (Distributed) Foundation Models/Agents.

Education

Nanyang Technological University, Singapore *2015 - 2020*
Ph.D, School of Computer Science and Engineering
Supervisor: Prof. Bo An

Dalian University of Technology *2011 - 2015*
Bachelor of Science, School of Physics and Optoelectronic Technology

Research Appointments

Assistant Professor *Jul. 2024 - Present*
School of Computing and Information Systems **Singapore Management University**

Research Assistant Professor *Oct. 2021 - Jun. 2024*
Agent Mediated Intelligence (Prof. Bo An) **Nanyang Technological University, Singapore**

Research Fellow *May. 2020 - Sep. 2021*
Agent Mediated Intelligence (Prof. Bo An) **Nanyang Technological University, Singapore**

Research Assistant *Aug. 2019 - Apr. 2020*
Agent Mediated Intelligence (Prof. Bo An) **Nanyang Technological University, Singapore**

Visiting Scholar *Jul. 2018 - Sep. 2018*
USC CAIS (Prof. Milind Tambe) **University of Southern California**

Selected Publications

(* indicates equal contribution, † indicates corresponding author)

[34]. Efficient integration of external knowledge to LLM-based world models via retrieval-augmented generation and reinforcement learning Chang Yang, Xinrun Wang†, Qinggang Zhang, Qi Jiang, Xiao Huang The 2025 Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2025

[33]. Qinggang Zhang, Zhishang Xiang, Yilin Xiao, Le Wang, Junhui Li, **Xinrun Wang**, Jinsong Su. FaithfulRAG: Fact-level conflict modeling for context-faithful retrieval-augmented generation. The 63rd Annual Meeting of the Association for Computational Linguistics (**ACL**), 2025. **SAC Highlights Award**

[32]. Aye Phyu Phyu Aung, Xinrun Wang†, Ruiyu Wang, Hau Chan, Bo An, Xiaoli Li, J. Senthilnath. Double oracle neural architecture search for game theoretic deep learning models. IEEE Transactions on Image Processing (**TIP**), 2025.

[31]. Weihao Tan, Wentao Zhang, Xinrun Xu, Haochong Xia, Ziluo Ding, Boyu Li, Bohan Zhou, Junpeng Yue, Jiechuan Jiang, Yewen Li, Ruyi An, Molei Qin, Chuqiao Zong, Longtao Zheng, YuJie Wu, Xiaoqiang Chai, Yifei Bi, Tianbao Xie, Pengjie Gu, Xiyun Li, Ceyao Zhang, Long Tian, Chaojie Wang, Xinrun Wang, Börje F. Karlsson, Bo An, Shuicheng YAN, Zongqing Lu. Cradle: Empowering foundation agents towards general computer control. Proceedings of the 42nd International Conference on Machine Learning (**ICML**), 2025.

[30]. Longtao Zheng, Zhiyuan Huang, Zhenghai Xue, **Xinrun Wang**, Bo An, Shuicheng Yan. AgentStudio: A toolkit for building general virtual agents. Proceedings of the 2025 International Conference on Learning Representations (**ICLR**), 2025.

[29]. Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, Longtao Zheng, **Xinrun Wang**†, Bo An†. A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist. Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data (**KDD**), 2024.

[28]. Pengdeng Li, Shuxin Li, Chang Yang, **Xinrun Wang**†, Shuyue Hu, Xiao Huang, Hau Chan, Bo An. Configurable mirror descent: Towards a unification of decision making. *Proceedings of the 41st International Conference on Machine Learning* (**ICML**), 2024.

[27]. **Xinrun Wang***, Chang Yang*, Shuxin Li, Pengdeng Li, Xiao Huang, Hau Chan and Bo An. Reinforcement Nash equilibrium solver. *Proceedings of the 33rd International Joint Conference on Artificial Intelligence* (**IJCAI**), 2024.

[26]. Pengdeng Li, Shuxin Li, Chang Yang, **Xinrun Wang**†, Xiao Huang, Hau Chan, Bo An. Self-adaptive PSRO: Towards an automatic population-based game solver. *Proceedings of the 33rd International Joint Conference on Artificial Intelligence* (**IJCAI**), 2024.

[25]. Longtao Zheng, Rundong Wang, **Xinrun Wang**†, Bo An†. Synapse: Trajectory-as-exemplar prompting with memory for computer control. *Proceedings of the 2024 International Conference on Learning Representations* (**ICLR**), 2024.

[24]. Weihao Tan, Wentao Zhang, Shanqi Liu, Longtao Zheng, **Xinrun Wang**†, Bo An†. True knowledge comes from practice: Aligning large language models with embodied environments via reinforcement learning. *Proceedings of the 2024 International Conference on Learning Representations* (**ICLR**), 2024.

[23]. Shanqi Liu, Dong Xing, Pengjie Gu, Bo An, Yong Liu, **Xinrun Wang**†. Greedy sequential execution: Solving homogeneous and heterogeneous cooperative tasks with a unified

framework. *Proceedings of the 2024 International Conference on Learning Representations* (**ICLR**), **Spotlight**, 2024.

[22]. Pengdeng Li*, Shuxin Li*, **Xinrun Wang**†, Jakub Cerny, Youzhi Zhang, Stephen Marcus McAleer, Hau Chan, Bo An. Grasper: A generalist pursuer for pursuit-evasion problems. *Proceeding of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems* (**AAMAS**), 2024.

[21]. Molei Qin, Shuo Sun, Wentao Zhang, Haochong Xia, **Xinrun Wang**†, Bo An†. EarnHFT: Efficient hierarchical reinforcement learning for high frequency trading. *Proceedings of the 38th AAAI Conference on Artificial Intelligence* (**AAAI**), 2024.

[20]. Haochong Xia, Shuo Sun, **Xinrun Wang**†, Bo An†. Market-GAN: Adding control to financial market data generation with semantic context. *Proceedings of the 38th AAAI Conference on Artificial Intelligence* (**AAAI**), 2024.

[19]. Pengdeng Li, Runsheng Yu, **Xinrun Wang**†, Bo An. Transition-informed reinforcement learning for large-scale Stackelberg mean-field games. *Proceedings of the 38th AAAI Conference on Artificial Intelligence* (**AAAI**), 2024.

[18]. Shuo Sun*, Molei Qin*, Wentao Zhang, Haochong Xia, Chuqiao Zong, Jie Ying, Yonggang Xie, Lingxuan Zhao, **Xinrun Wang**†, Bo An†. TradeMaster: A holistic quantitative trading platform empowered by reinforcement learning. *Thirty-seventh Conference on Neural Information Processing Systems* (**NeurIPS**) Dataset and Benchmark Track, 2023.

[17]. Pengjie Gu, Xinyu Cai, Dong Xing, **Xinrun Wang**†, Mengchen Zhao, Bo An. Offline RL with discrete proxy representations for generalizability in POMDPs. *Thirty-seventh Conference on Neural Information Processing Systems* (**NeurIPS**), 2023.

[16]. Shuo Sun, **Xinrun Wang**†, Wanqi Xue, Xiaoxuan Lou, Bo An†. Mastering stock markets with efficient mixture of diversified trading experts. *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data* (**KDD**), 2023.

[15]. Shuo Sun, Molei Qin, **Xinrun Wang**†, Bo An†. PRUDEX-Compass: Towards systematic evaluation of reinforcement learning in financial markets. *Transactions on Machine Learning Research* (**TMLR**), 2023

[14]. Runsheng Yu, Weiyu Chen, **Xinrun Wang**, James Kwok. Enhancing meta learning via multi-objective soft improvement functions. *The 11th International Conference on Learning Representations* (**ICLR**), 2023.

[13]. Pengdeng Li, **Xinrun Wang**†, Shuxin Li, Hau Chan, Bo An. Population-size-Aware Policy Optimization for Mean-Field Games. *The 11th International Conference on Learning Representations* (**ICLR**), 2023.

[12]. Shuxin Li, **Xinrun Wang**†, Youzhi Zhang†, Wanqi Xue, Jakub Cerny, Bo An. Solving large-scale pursuit-evasion games using pre-trained strategies. *Proceedings of the 37th AAAI Conference on Artificial Intelligence* (**AAAI**), 11586-11594, 2023.

[11]. Aye Phyu Phyu Aung*, **Xinrun Wang***†, Runsheng Yu*, Bo An†, Senthilnath Jayavelu, Xiaoli Li†. DO-GAN: A double oracle framework for generative adversarial networks. *Proceedings of the 2022 IEEE Conference on Computer Vision and Pattern Recognition* (**CVPR**), pp. 11275-11284, 2022.

[10]. Wei Qiu, **Xinrun Wang**, Runsheng Yu, Rundong Wang, Xu He, Bo An, Svetlana Obraztsova, Zinovi Rabinovich. RMIX: Learning risk-sensitive policies for cooperative reinforcement learning agents. *Proceedings of the 35th Annual Conference on Neural Information Processing Systems* (**NeurIPS**), pp. 23049-23062, 2021.

[9]. Yanchen Deng, Runsheng Yu, **Xinrun Wang**, Bo An. Neural regret matching for distributed constraint optimization problems. *Proceedings of the 30th International Joint Conference on Artificial Intelligence* (**IJCAI**), pp. 146-153, 2021.

[8]. Shuxin Li, Youzhi Zhang*, **Xinrun Wang***, Wanqi Xue, Bo An. CFR-MIX: Solving imperfect information extensive-form games with combinatorial action space. *Proceedings of the 30th International Joint Conference on Artificial Intelligence* (**IJCAI**), pp. 3663-3669, 2021.

[7]. Wanqi Xue, Youzhi Zhang, Shuxin Li, **Xinrun Wang**, Bo An, Chai Kiat Yeo. Solving large-scale extensive-form network security games via neural fictitious self-play. *Proceedings of the 30th International Joint Conference on Artificial Intelligence* (**IJCAI**), pp. 3713-3720, 2021.

[6]. Aye Phyu Phyu Aung, **Xinrun Wang**, Bo An, Xiaoli Li. We mind your well-being: Preventing depression in uncertain social networks by sequential interventions. *Proceedings of the 30th International Conference on Automated Planning and Scheduling* (**ICAPS**), pp. 499-507, 2020.

[5]. Zhenyu Shi*, Runsheng Yu*, **Xinrun Wang***, Rundong Wang, Youzhi Zhang, Hanjiang Lai, Bo An. Learning expensive coordination: An event-based deep RL approach. *Proceedings of the 2020 International Conference on Learning Representations* (**ICLR**), 2020.

[4]. **Xinrun Wang**, Milind Tambe, Branislav Bosansky, Bo An. When players affect target values: Modeling and solving dynamic partially observable security games. *Proceedings of the 10th Conference on Decision and Game Theory for Security* (**GameSec**), pp. 542-562, 2019. (**Outstanding Student Paper Award**)

[3]. **Xinrun Wang**, Bo An, Hau Chan. Who should pay the cost: A game-theoretic model for government subsidized investments to improve national cybersecurity. *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (**IJCAI**), pp. 6020-6027, 2019.

[2]. **Xinrun Wang**, Bo An, Martin Strobel, Fookwai Kong. Catching Captain Jack: Efficient time and space dependent patrols to combat oil-siphoning in international waters. *Proceedings of the 32nd AAAI Conference on Artificial Intelligence* (**AAAI**), pp. 208-215, 2018.

[1]. **Xinrun Wang**, Qingyu Guo, Bo An. Stop nuclear smuggling through efficient container inspection. *Proceedings of the 16th International Joint Conference on Autonomous Agents and Multi-Agent Systems* (**AAMAS**), pp. 669-677, 2017.

Open-Sourced Projects

TradeMaster: A first-of-its kind, best-in-class open-source platform for quantitative trading (QT) empowered by reinforcement learning (RL), which covers the full pipeline for the design, implementation, evaluation and deployment of RL-based algorithms.
Project website: https://github.com/TradeMaster-NTU/TradeMaster

Cradle: A novel modular Large Language Model (LMM)-powered framework that empowers foundation agents towards **general computer control**, i.e., Building foundation agents that can master **ANY** computer task via the universal human-style interface by receiving input from screen and audio and outputting keyboard and mouse actions..
Project website: https://baai-agents.github.io/Cradle/

Awards

Lee Kong Chian Fellow *2025-2027*

Outstanding Student Paper Award *GameSec, 2019*

National Scholarship *China, 2013*

National Encouragement Scholarship *China, 2012*

Professional Services

Local Chair: DAI 2024

Workshop Chair: OptLearnMAS (2024, 2025)

Area Chair: ICLR 2026, NeurIPS 2025, AAMAS (2023, 2024)

Conference Program Committee Member (Reviewer):
ICML (2021,2022); NeurIPS (2020,2021,2022,2023); ICLR (2022,2023);
AAAI (2021,2022, 2023); IJCAI (2021,2022,2023); AAMAS (2022)

Journal Reviewer:
TNNLS, JAAMAS, JAIR