ZIYAN WANG

Phone: (+44) 742-190-5361 ◇ Email: ziyan.wang@kcl.ac.uk
Homepage: ziyanwang98.github.io
Google Scholar ◇ Github ◇ LinkedIn

EDUCATION

King's College London (KCL) Dec. 2022 - Now
Ph.D. in Computer Science London, UK

· Advisors: Prof. Yali Du, Prof. Sanjay Modgil

Carnegie Mellon University (CMU) Spring 2025 (Expected)
Visiting Ph.D. Student Pittsburgh, US

· Advisor: Prof. Fei Fang

University College London (UCL) Sep. 2020 - Dec. 2021
M.S. in Robotics and Computation London, UK

· Grade: 72/100 (Distinction) Rank: top 15%
· Advisors: Prof. Jun Wang

Eötvös Loránd University (ELTE) Sep. 2016 - Jun. 2019
B.S. in Computer Science Budapest, Hungary

· GPA: 4.73/5.0 Rank: 3/78

RESEARCH INTERESTS

My research interests focus on Multi-agent Reinforcement Learning (**MARL**), Large Language Models (**LLMs**), and **Robotics**. I am particularly interested in developing algorithms for agent collaboration in complex environments, enabling effective human-robot communication through natural language understanding, ensuring safe policy learning under constraints, and exploring LLMs' potential in multi-agent scenarios.

PUBLICATIONS

* denotes equal contribution.

Refereed Conference Publications

1. Policy Learning from Tutorial Books via Understanding, Rehearsing and Introspecting
Xiong-Hui Chen*, **Ziyan Wang***, Yali Du, Shengyi Jiang, Meng Fang, Yang Yu, and Jun Wang
The Conference on Neural Information Processing Systems (NeurIPS 2024 **Oral**)

2. Learning to Discuss Strategically: A Case Study on One Night Ultimate Werewolf
Xuanfa Jin*, **Ziyan Wang***, Yali Du, Meng Fang, Haifeng Zhang, and Jun Wang
The Conference on Neural Information Processing Systems (NeurIPS 2024)

3. Chessgpt: Bridging policy learning and language modeling
Xidong Feng, Yicheng Luo, **Ziyan Wang**, et.al
The Conference on Neural Information Processing Systems (NeurIPS 2023)

4. Interpretable Reward Redistribution in Reinforcement Learning: A Causal Approach
   Yudi Zhang, Yali Du, Biwei Huang, **Ziyan Wang**, et.al
   The Conference on Neural Information Processing Systems (NeurIPS 2023)

5. Sauté RL: Almost Surely Safe Reinforcement Learning Using State Augmentation
   Aivar Sootla, Alexander I. Cowen-Rivers, Taher Jafferjee, **Ziyan Wang**, et.al
   The International Conference on Machine Learning (ICML 2022)

6. Safe Reinforcement Learning with Free-form Natural Language Constraints and Pre-Trained Language Models
   Xingzhou Lou, Junge Zhang, **Ziyan Wang**, et.al
   The International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2024)

Conference Workshop Publications

1. Safe Multi-agent Reinforcement Learning with Natural Language Constraints
   Ziyan Wang, Meng Fang, Tristan Tomilin, Fei Fang, and Yali Du
   ICLR 2024 Workshop on Generative Models for Decision Making (ICLR GenAI4DM)

2. MACCA: Offline Multi-agent Reinforcement Learning with Causal Credit Assignment
   Ziyan Wang, Yali Du, Yudi Zhang, Meng Fang, and Biwei Huang
   In NeurIPS 2024 Causal Representation Learning Workshop (NeurIPS CRL)

Preprints

1. M^{3}HF: Multi-agent Reinforcement Learning from Multi-phase Human Feedback of Mixed Quality
   Ziyan Wang, Zhicheng Zhang, Fei Fang, Yali Du
   Under Review

2. Multi-agent constrained policy optimisation
   Shangding Gu, Jakub Grudzien Kuba, Munning Wen, Ruiqing Chen, **Ziyan Wang**, et.al
   Under Review

ACHIEVEMENTS

NeurIPS 2024 Scholar Award	*Oct. 2024*
NeurIPS 2024 Oral Presentation	*Sep. 2024*
King's PhD Scholarship	*Dec. 2022 - Dec.2026*
Hungarian State Scholarship	*Sep. 2016 - Jun.2019*

TEACHING

Teaching, Oxford Machine Learning Summer School - DNN + Optimisation	*Summer 2024*
Teaching Assistant, Oxford Machine Learning Summer School (ML x Fundamentals)	*Summer 2023*
Teaching Assistant, Optimisation Methods (Level 6), King's College London	*Spring 2023*

ACADEMIC SERVICE

Conference Reviewer

International Conference on Machine Learning (ICML)	*2023/2024*
Conference on Neural Information Processing Systems (NeurIPS)	*2023/2024*
International Conference on Learning Representations (ICLR)	*2024/2025*

International Conference on Artificial Intelligence and Statistics (AISTATS) *2025*

Journal Reviewer

IEEE Transactions on Knowledge and Data Engineering (TKDE) *2023-*

IEEE Transactions on Artificial Intelligence (TAI) *2023-*

SKILLS

Programming Languages	Python, C/C++, Java, MATLAB, LaTeX
Learning Frameworks	PyTorch, Jax, Numpy, Pandas
Robotics	MuJoCo, IsaacGym, OpenAI Gym
Others	Ray, RLlib, Git, Linux, Docker