Zhicheng Zhang

zhichen3@cs.cmu.edu | gabr1e1.github.io | +1 (412) 996-5392 | TCS Hall 463, 4665 Forbes Ave Last updated: Dec 2023

Education

School of Computer Science, Carnegie Mellon University Sep 2022 - Present
- Ph.D. Student, Societal Computing Pittsburgh, PA
- Advisor: Fei Fang

ACM Honors Class, Shanghai Jiao Tong University Sep 2018 - Jun 2022
- B.Eng. in Computer Science, GPA: 3.98/4.3 Shanghai, China
- Scholarships: National Scholarship 2019 (top 0.2% national-wide), Foresight Foundation-Sequoia Capital Talent Development Scholarship 2021
- Advisors: Weinan Zhang & Yong Yu

Research Interests

Explainable Reinforcement Learning (XRL), (Model-based) Multi-agent RL, Explainable AI, AI for public health

Research Experience

Carnegie Mellon University, Graduate Research Assistant Sep 2022 - Present

AI and Social Good Lab (AISOC Lab), advised by Fei Fang Pittsburgh, PA

- **Natural-Language Explanation Generation for Chess AI (ongoing)**.
- **User-Centered Approach to Interpretable Decision-Tree Policies for Reinforcement Learning (ongoing)**: Developed a user-centered approach for explainable reinforcement learning, utilizing simulated user feedback and end-use cases in a feedback loop to enhance the training of interpretable decision-tree policies.
- **Learning Decision Tree Policies for Interpretable Multi-Agent Reinforcement Learning**. Proposed a novel centralized decision-tree training method for extracting interpretable decision-tree policies from neural networks policies trained with multi-agent reinforcement learning algorithms.
- **Explanation Generation for Task Difficulty Prediction in Crowdsourcing Food Rescue Platforms**.Developed three user-friendly explanation methods utilizing Large Language Models (LLMs) to generate interpretable insights for volunteers to better understand the task difficulty predictions. The models and findings from the complete study are in the process of being adopted at Food Rescue Hero, a large food rescue platform serving over 25 cities across the United States.
- **Pandemic Intelligence through Wastewater-Based Epidemiology (WBE)**. Demonstrated the efficacy of using only publicly available wastewater and case count data to forecast COVID-19 spread, implementing an innovative segmentation-based evaluation metric and addressing data shortcomings, which enables forecasting accuracy comparable to methods with more extensive data resources.

Carnegie Mellon University, Research Intern May 2021 - Jan 2022

AI and Social Good Lab (AISOC Lab), advised by Fei Fang Remote

- **Multi-Agent Meta-Exploration through Exploiting State-Action Space Structure**. Designed a novel meta-exploration method for cooperative multi-agent learning in sparse-reward environments that first identifies the agents' high-rewarding joint state-action subspace from meta-training tasks and then learns a set of diverse exploration policies to "cover" the subspace.

Shanghai Jiao Tong University, Undergraduate Researcher Jul 2020 - Apr 2022

Apex Lab, advised by Weinan Zhang & Yong Yu Shanghai, China

- **Model-based Offline Policy Optimization with Distribution Correcting Regularization**. Established a theoretical lower bound on the return in model-based offline RL with an explicit account of model uncertainty measured by model error. Estimated the density ratio between model rollout distribution and offline dataset distribution via the DICE framework to regularize model-predicted rewards.
- **Local-Global Model-based Approach for City-scale Traffic Light Control**, KDD Cup 2021 City Brain Challenge, *Finalist (Top 20/1156 teams)*. Built high-precision car flow dynamics model using real-world data, incorporating varying congestion levels and accounting for time in lane-changing behaviors. Achieved ~205% increase in total served vehicles compared to traditional fixed-interval control methods under real city-scale traffic data.

Work Experience

Metabit Trading (Qianxiang Investment) May 2022 - Aug 2022
Quantitative Research Intern Shanghai, China

- Integrated from scratch Level 2 market data for China futures with parallelized and scalable data integrity checks and evaluation pipeline
- Researched trading signals with improved inference speed using Numba and higher prediction accuracy

Publications

Preprints

- X. Wang, **Z. Zhang**, W. Zhang. Model-based Multi-agent Reinforcement Learning: Recent Progress and Prospects. arXiv preprint arXiv:2203.10603, 2022.

Book Chapters

- S. Milani, **Z. Zhang**, N. Topin, Z. R. Shi, C. Kamhoua, E. E. Papalexakis, and F. Fang. Interpretable Multi-Agent Reinforcement Learning with Decision-Tree Policies. To appear in Explainable Agency in Artificial Intelligence, CRC Press/Taylor&Francis, 2024.

Peer-Reviewed Conference Papers and Proceedings

- Z. R. Shi, J. Zhi, S. Zeng, **Z. Zhang**, A. Kapoor, S. Hudson, H. Shen and F. Fang. Predicting and Presenting Task Difficulty for Crowdsourcing Food Rescue Platforms. WWW 2024 - Web4Good Track [In submission].
- **Z. Zhang***, Y. Liang*, Y. Wu, and F. Fang. MESA: Multi-Agent Meta-Exploration through Exploiting State-Action Space Structure. AAMAS 2024.
- S. Milani*, **Z. Zhang***, N. Topin, Z. R. Shi, C. Kamhoua, E. E. Papalexakis, and F. Fang. MAVIPER: Learning Decision Tree Policies for Interpretable Multi-Agent Reinforcement Learning. ECML 2022.
- J. Shen*, M. Chen*, **Z. Zhang**, Z. Yang, W. Zhang, and Y. Yu. Model-based Offline Policy Optimization with Distribution Correcting Regularization. ECML 2021.
- M. Zhu*, M. Liu*, J. Shen, **Z. Zhang**, S. Chen, W. Zhang, D. Ye, Y. Yu, Q. Fu, and W. Yang. MapGo: Model-Assisted Policy Optimization for Goal-Oriented Tasks. IJCAI 2021.

Peer-Reviewed Workshop Papers and Extended Abstracts

- **Z. Zhang**, S. Neumeister, A. Desai, M. S. Majumder, and F. Fang. Unlocking the Potential of Public Datasets: Wastewater-Based Epidemiological Forecasting During COVID-19. In epiDAMIK 6.0 workshop at KDD 2023.

Professional Services

Predictive Intelligence for Pandemic Prevention (PIPP) Initiative

- Supported by the National Science Foundation (NSF)
- Co-organized the *Modeling Intervention Acceptance for Disease Mitigation Workshop* Apr 2023
- Moderated the panel "Co-evolving Outbreaks: Understanding the Interdependence of Disease and (Mis)Information Spread" for the *Predictive Intelligence for Limiting Outbreak Threats (PILOT) Synthesis Workshop* Sep 2023

GameSec 2022 Conference, Reviewer 2022
ICML 2022, Reviewer 2022

Talks

CHIP Fellows' Symposium Mar 2023

- Gave a 10-minute talk on *Interpretable Multi-Agent Reinforcement Learning* at the Boston Children's Hospital Computational Health Informatics Program (CHIP) Fellows' Symposium

Teaching

Shanghai Jiao Tong University, Teaching Assistant

- Data Structure (CS147) Spring 2020
- Practice of Computer Algorithms (MS125) Summer 2020

Carnegie Mellon University, Teaching Assistant

- AI Methods for Social Good (17-737) Spring 2024

Skills

Programming Languages: Python, C, C++, Java, Verilog HDL
Natural Languages: Mandarin Chinese (native), English (fluent), Spanish (proficient)