



Shun Zhang

✉ shunzh@umich.edu  [shunzh.github.io](https://github.com/shunzh)  San Francisco Bay Area

Research interests: Reinforcement learning; large language models; automated theorem proving; automatic code generation; value alignment.

Experience

Research Scientist, *MIT-IBM Watson AI Lab* Jun. 2022 - Present
Postdoctoral Researcher, *MIT-IBM Watson AI Lab* Oct. 2021 - Jun. 2022
Postdoctoral Researcher, *IBM-NJIT* Aug. 2020 - Oct. 2021

- Research and publish papers on **reinforcement learning** and **large language models**, with a focus on competitive-level code generation, reinforcement learning from human feedback, and AI for electric circuit design automation.

Graduate Research Assistant, *University of Michigan* (Ann Arbor, MI) Sep. 2015 - Apr. 2020

- Conducted research and published papers on **preference elicitation** and **AI safety in reinforcement learning**.
- Designed active learning algorithms to improve a learning agent's performance and guarantee safety in domains with uncertain objectives.

Software Development Engineer Intern, *Amazon* (Seattle, WA) Jun. - Aug. 2014

- Created a WebRTC-related internal tool to resolve cross-departmental communication issues.

Software Development Engineer Intern, *Semantic Designs* (Austin, TX) Jun. - Aug. 2013

- Created a user interface for a programming language analysis tool for better visualization.

Education

Ph.D. in Computer Science and Engineering, *University of Michigan* Sep. 2015 - Apr. 2020

- Dissertation: Efficiently Finding Approximately-Optimal Queries for Improving Policies and Guaranteeing Safety
- Advisors: Satinder Singh, Edmund H. Durfee

M.S. in Computer Science, *University of Texas at Austin* Aug. 2015

- Master thesis: Parameterized Modular Inverse Reinforcement Learning
- Committee members: Dana Ballard, Peter Stone

B.S. in Computer Science, *University of Texas at Austin* May 2014

Publications and Preprints

- **Improving Reinforcement Learning from Human Feedback with Efficient Reward Model Ensemble (Short Paper)**
Shun Zhang, Zhenfang Chen, Sunli Chen, Yikang Shen, Zhiqing Sun, and Chuang Gan
arXiv, 2024
[paper](#)
- **Adaptive Online Replanning with Diffusion Models**
Siyuan Zhou, Yilun Du, Shun Zhang, Mengdi Xu, Yikang Shen, Wei Xiao, Dit-Yan Yeung, and Chuang Gan
Conference on Neural Information Processing Systems (NeurIPS), 2023
[paper](#)

- **Planning with Large Language Models for Code Generation**
Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B. Tenenbaum, and Chuang Gan
International Conference on Learning Representations (ICLR), 2023
[paper](#)
- **Hyper-Decision Transformer for Efficient Online Policy Adaptation**
Mengdi Xu, Yuchen Lu, Yikang Shen, Shun Zhang, Ding Zhao, and Chuang Gan
International Conference on Learning Representations (ICLR), 2023
[paper](#)
- **Prompting Decision Transformer for Few-shot Policy Generalization**
Mengdi Xu, Yikang Shen, Shun Zhang, Yuchen Lu, Ding Zhao, Joshua B. Tenenbaum, and Chuang Gan
International Conference on Machine Learning (ICML), 2022
[paper](#)
- **Power Converter Circuit Design Automation using Parallel Monte Carlo Tree Search**
Shaoze Fan, Shun Zhang, Jianbo Liu, Ningyuan Cao, Xiaoxiao Guo, Jing Li, and Xin Zhang
ACM Transactions on Design Automation of Electronic Systems (TODAES), 2022
[paper](#)
- **From Specification to Topology: Automatic Power Converter Design via Reinforcement Learning**
Shaoze Fan, Ningyuan Cao, Shun Zhang, Jing Li, Xiaoxiao Guo, and Xin Zhang
International Conference on Computer Aided Design (ICCAD), 2021
[paper](#)
- **Efficiently Finding Approximately-Optimal Queries for Improving Policies and Guaranteeing Safety**
Shun Zhang
Ph.D. Dissertation, 2020
[paper](#)
- **Querying to Find a Safe Policy Under Uncertain Safety Constraints in Markov Decision Processes**
Shun Zhang, Edmund H. Durfee, and Satinder Singh
AAAI Conference on Artificial Intelligence (AAAI), 2020
[paper](#)
- **Minimax-Regret Querying on Side Effects for Safe Optimality in Factored Markov Decision Processes**
Shun Zhang, Edmund H. Durfee, and Satinder Singh
International Joint Conference on Artificial Intelligence (IJCAI), 2018
[paper](#)
- **Modeling Sensory-Motor Decisions in Natural Behavior**
Ruohan Zhang, Shun Zhang, Matthew H. Tong, Yuchen Cui, Constatin A. Rothkopf, Dana H. Ballard, and Mary M. Hayhoe
PLoS Computational Biology, 2018
[paper](#)
- **Approximately-Optimal Queries for Planning in Reward-Uncertain Markov Decision Processes**
Shun Zhang, Edmund H. Durfee, and Satinder Singh
International Conference on Automated Planning and Scheduling (ICAPS), 2017
[paper](#)
- **Determining Placements of Influencing Agents in a Flock**
Katie Genter, Shun Zhang, and Peter Stone
Autonomous Agents and Multiagent Systems (AAMAS), 2015
[paper](#)

- **Autonomous Intersection Management for Semi-Autonomous Vehicles**
Tsz-Chiu Au, **Shun Zhang**, and Peter Stone
Handbook of Transportation, 2015
[paper](#)

Academic Services

Conference Reviewer

IEEE ITSC 2014, AAAI 2019, AISTATS 2023-24, CVPR 2023, ICML 2023-24, NeurIPS 2023, ICLR 2024.

Skills

Research

Reinforcement learning, convex optimization, deep learning, large language models, active learning, planning under uncertainty.

Programming languages

Proficient in Python (NumPy, PyTorch). Experienced in Java, C++, C, Scheme, Matlab.

Miscellaneous

- Languages: Mandarin Chinese (native), English (professional proficiency).
- No sponsorship required to work in the United States.