



Shun Zhang

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

Research interests: Reinforcement learning; large language models; code generation and reasoning; value alignment; artificial general intelligence.

Experience

- Senior Generative AI Engineer, NVIDIA (Santa Clara, CA)** Aug 2025 - Present
- Training and evaluating vision-language models (VLMs) and world action models (WAMs) for **physical AI**.
- Founding Member of Technical Staff, Asari AI (San Francisco, CA)** Jun 2024 - Jan 2025
- Developed a **code agent** that plans, verifies, and discovers new skills and knowledge.
- Research Scientist, MIT-IBM Watson AI Lab** Jun 2022 - Jun 2024
- Research on **reinforcement learning** and post-training of **language models**, with a focus on code generation and reinforcement learning from human feedback.
- Postdoctoral Researcher, IBM Research** Aug 2020 - Jun 2022
- Research on **meta-reinforcement learning** and **AI for scientific discovery**.
- Software Development Engineer Intern, Amazon (Seattle, WA)** Jun 2014 - Aug 2014
- Created a WebRTC-related internal tool to resolve cross-departmental communication issues.
- Software Development Engineer Intern, Semantic Designs (Austin, TX)** Jun 2013 - 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 (Ann Arbor, MI)** Sep 2015 - Apr 2020
- Advisors: [Prof. Satinder Singh](#) and [Prof. Ed Durfee](#).
 - Research on value alignment and AI safety in reinforcement learning.
- B.S. and M.S. in Computer Science, University of Texas at Austin (Austin, TX)** Aug 2015
- Undergraduate/master research advisors: [Prof. Peter Stone](#) and [Prof. Dana Ballard](#).

Publications and Preprints

- **Cosmos 3: Omnimodal World Models for Physical AI**
NVIDIA
arXiv, 2026
[paper](#)
- **Improving Reinforcement Learning from Human Feedback with Efficient Reward Model Ensemble**
Shun Zhang, Zhenfang Chen, Sunli Chen, Yikang Shen, Zhiqing Sun, and Chuang Gan
arXiv, 2024
[paper](#)
- **Graph-Transformer-based Surrogate Model for Accelerated Converter Circuit Topology Design**
Shaoze Fan, Haoshu Lu, **Shun Zhang**, Ningyuan Cao, Xin Zhang, and Jing Li
Design Automation Conference (DAC), 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

AAAI 2019, AISTATS 2023-24, CVPR 2023, ICML 2023-24, NeurIPS 2023-25, ICLR 2024-25.