

Meng Song

🏠 <https://mengsong16.github.io/> ✉ mes050@ucsd.edu ☎ +1(412) 897-7893

RESEARCH INTERESTS

Reinforcement Learning, Self-Supervised Learning

EDUCATION

University of California, San Diego 2016 - 2024
Ph.D. in Computer Science

Carnegie Mellon University 2013 - 2014
M.S. in Robotics

WORK EXPERIENCE

Pinetree Research, Member of Technical Staff October 2025 – Present

- Independently developed a browser agent with autonomous reflection-guided retries. Achieved 91% task success on Online-Mind2Web in internal human evaluation with up to three attempts per task, compared with the then-listed human-evaluation leaderboard best of 78.7% (Feb. 2026).
- Independently developed an end-to-end voice-agent system with response precomputation from partial transcripts, input-consistency checks, and streaming speech synthesis. Optimized the pipeline to balance latency and speech quality, achieving comparable latency and conversational reliability to an ElevenLabs Flash v2 baseline in internal multi-turn evaluations.
- Investigate how LLM agents select goal-relevant evidence and bind constraints to specific entities to guide downstream decisions, using targeted KV-cache interventions. Preliminary enterprise-workflow experiments showed goal-dependent changes in eligibility judgments and recommendations, while unrelated judgments remained unchanged.

Handshake AI Solutions, MOVE Fellow April 2025 – October 2025

- Contributed to a confidential research project for a leading frontier AI lab, focused on improving LLMs' ability to conduct machine learning research.

Snap, Research Scientist Intern Summer 2017

- Developed a contrastive learning-based logo detector conditioned on a single reference image, trained on 130K+ annotated images. Improved recall over classification baselines by 41 percentage points on unseen classes and 59 on seen classes.

SELECTED PROJECTS

University of California, San Diego, Graduate Student Researcher

Probabilistic World Modeling 2023–2024

- Independently developed a probabilistic world model using asymmetric contrastive learning to encode probabilistic state reachability. Enabled planning via latent-space midpoint interpolation and identified bottleneck subgoals through clustering to support structure-aware exploration and long-horizon task decomposition.

Zero-Shot Policy Learning with Decision Transformers 2023

- Independently developed a Decision Transformer prompting method combining task parameters with jointly learned prompt embeddings, enabling zero-shot generalization without test-time demonstration prompts. Improved normalized scores over a demonstration-conditioned baseline on multiple control, manipulation, and navigation tasks, with gains of up to 25%.

Reward Design for RL-Based Prompt Optimization

2022

- Designed a piecewise, margin-based reward for RLPrompt to align discrete prompt optimization with classification accuracy. Ablation studies showed improved training stability and average classification accuracy across random seeds.

RL and Imitation Learning for Visual Navigation

2021–2022

- Independently implemented and trained recurrent PPO and behavior cloning agents in Habitat with parallel rollout collection. Showed that PPO retained higher zero-shot success rates even when BC received additional expert demonstrations. Proposed TD-based stitching of suboptimal and failed trajectories as an explanation for PPO’s stronger generalization beyond imitation of expert behavior.

Embodied Simulation and PPO Training

2020–2021

- Independently developed an embodied simulation environment using OpenRooms assets and Bullet physics, with configurable materials, object masses, and surface friction. Designed two rearrangement tasks and physics-informed rewards to encourage physics-aware behavior in PPO agents. Learned policies favored lighter objects and lower-friction paths, reducing physical cost while maintaining comparable task success rates.

SELECTED PUBLICATIONS

- **Meng Song**, “Good Actions Succeed, Bad Actions Generalize: A Case Study on Why RL Generalizes Better”, in *Out-of-Distribution Generalization in Robotics Workshop at The Robotics: Science and Systems* (RSS) [Spotlight], 2025 [PDF]
- **Meng Song**, “Probabilistic World Modeling with Asymmetric Distance Measure”, in *Geometry-grounded Representation Learning and Generative Modeling Workshop at International Conference on Machine Learning* (ICML) [Oral], 2024 [PDF]
- **Meng Song**, Xuezhi Wang, Tanay Biradar, Yao Qin, Manmohan Chandraker, “A Minimalist Prompt for Zero-Shot Policy Learning”, in *Task Specification Workshop at The Robotics: Science and Systems* (RSS), 2024 [PDF]
- Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, **Meng Song**, Eric P. Xing, Zhiting Hu, “RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning”, in *Conference on Empirical Methods in Natural Language Processing* (EMNLP), 2022 [PDF]
- **Meng Song**, Yuhan Liu, Zhengqin Li, Manmohan Chandraker, “Learning to Rearrange with Physics-Inspired Risk Awareness”, in *Risk Aware Decision Making Workshop at The Robotics: Science and Systems* (RSS), 2022 [PDF]
- Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, **Meng Song**, Yuhan Liu, et al., “OpenRooms: An End-to-End Open Framework for Photorealistic Indoor Scene Datasets”, in *Conference on Computer Vision and Pattern Recognition* (CVPR) [Oral], 2021 [PDF]

SKILLS

Programming & ML: Python, C++, PyTorch

Reinforcement Learning: PPO, SAC, reward design, parallel rollout collection

Simulation Environments: Habitat, MuJoCo, PyBullet