

## EDUCATION

---

<b>University of California San Diego</b> Ph.D. in Computer Science and Engineering, <b>Advisor:</b> Prof. Hao Su	San Diego, USA 2019 - Current
<b>University of California San Diego</b> M.S. in Computer Science and Engineering, <b>GPA:</b> 4.0/4.0	San Diego, USA 2017 - 2019
<b>Zhejiang University</b> B.Eng. in Computer Science and Technology, <b>GPA:</b> 3.8/4.0, <b>Major GPA:</b> 3.97/4.00	Hangzhou, China 2013 - 2017

## RESEARCH INTERESTS

---

I study **embodied AI** from a **data-centric** perspective. My previous research can be summarized as “**modeling the world for AI agents**” (see [this figure](#) for a brief overview). More specifically, my focus areas include **Foundation Models for Robotics**, **Scaling Up Robotic Data Collection**, as well as **Reinforcement Learning and Imitation Learning**.

## INDUSTRY EXPERIENCES

---

<b>1X Technologies</b> Robot Learning Intern – Project: Multimodal Large Language Models (MLLMs) for Skill Chaining	Sunnyvale, United States Summer 2024
<b>NVIDIA Seattle Robotics Lab</b> Research Intern – Project: Data-Efficient Demonstration Collection for Robotic Foundation Models	Seattle, United States Summer 2023
<b>Amazon Alexa AI</b> Applied Scientist Intern – Project: Interpretable RL for Text-Based Games with Graph Inputs	Sunnyvale, United States Summer 2021
<b>Wormpex AI Research</b> Research Intern – Project: Store Layout Optimization based on Customer Behavior Model	Remote, United States Summer 2020
<b>Intel AI</b> Research Intern – Project: Memory-Constrained Navigation via Combining RL and Planning	San Diego, United States Summer 2019
<b>Microsoft Research Asia</b> Research Intern – Project: Indoor Visual Navigation by Deep RL	Beijing, China Apr 2017 - Aug 2017

## UNDER REVIEW SUBMISSIONS

---

1. Xiu Yuan\*, **Tongzhou Mu\***, Stone Tao, Yunaho Fang, Michael Zhang, Hao Su, “Policy Decorator: Model-Agnostic Online Refinement for Large Policy Model”, *Under Review*, 2024
2. **Tongzhou Mu\***, Zhaoyang Li\*, Stanisław Wiktor Strzelecki\*, Xiu Yuan, Yunchao Yao, Litian Liang, Hao Su, “When Should We Prefer State-to-Visual DAgger Over Visual Reinforcement Learning?”, *Under Review*, 2024

## PUBLICATIONS & PREPRINTS

---

\* and † indicate equal contributions.

- [1] S. Tao, F. Xiang, A. Shukla, Y. Qin, X. Hinrichsen, X. Yuan, C. Bao, X. Lin, Y. Liu, T.-k. Chan, Y. Gao, X. Li, **T. Mu**, N. Xiao, A. Gurha, Z. Huang, R. Calandra, R. Chen, S. Luo, and H. Su, “Maniskill3: GPU parallelized robotics simulation and rendering for generalizable embodied AI”, *arXiv preprint*, 2024.
- [2] Z. Ling, Y. Fang, X. Li, **T. Mu**, M. Lee, R. Pourreza, R. Memisevic, and H. Su, “Unleashing the creative mind: Language model as hierarchical policy for improved exploration on challenging problem solving”, *arXiv preprint*, 2024.
- [3] **T. Mu**, Y. Guo, J. Xu, A. Goyal, H. Su, D. Fox, and A. Garg, “AdaDemo: Data-efficient demonstration expansion for generalist robotic agent”, in *International Symposium of Robotics Research (ISRR)*, 2024.
- [4] **T. Mu**, M. Liu, and H. Su, “DrS: Learning reusable dense rewards for multi-stage tasks”, in *International Conference on Learning Representations (ICLR)*, 2024.
- [5] N. Hansen\*, Z. Yuan\*, Y. Ze\*, **T. Mu\***, A. Rajeswaran, H. Su, H. Xu, and X. Wang, “On pre-training for visuo-motor control: Revisiting a learning-from-scratch baseline”, in *International Conference on Machine Learning (ICML)*, 2023.
- [6] S. Tao, X. Li, **T. Mu**, Z. Huang, Y. Qin, and H. Su, “Abstract-to-executable trajectory translation for one-shot task generalization”, in *International Conference on Machine Learning (ICML)*, 2023.
- [7] **T. Mu** and H. Su, “Boosting reinforcement learning and planning with demonstrations: A survey”, *arXiv preprint*, 2023.
- [8] J. Gu†, F. Xiang†, X. Li\*, Z. Ling\*, X. Liu\*, **T. Mu\***, Y. Tang\*, S. Tao\*, X. Wei\*, Y. Yao\*, X. Yuan, P. Xie, Z. Huang, R. Chen, and H. Su, “ManiSkill2: A unified benchmark for generalizable manipulation skills”, in *International Conference on Learning Representations (ICLR)*, 2023.
- [9] X. Zhang, R. Chen, A. Li, F. Xiang, Y. Qin, J. Gu, Z. Ling, M. Liu, P. Zeng, S. Han, Z. Huang, **T. Mu**, J. Xu, and H. Su, “Close the optical sensing domain gap by physics-grounded active stereo sensor simulation”, *IEEE Transactions on Robotics (T-RO)*, pp. 1–19, 2023.
- [10] **T. Mu**, K. Lin, F. Niu, and G. Thattai, “Learning two-step hybrid policy for graph-based interpretable reinforcement learning”, *Transactions on Machine Learning Research (TMLR)*, 2022.
- [11] **T. Mu\***, Z. Ling\*, F. Xiang\*, D. C. Yang\*, X. Li\*, S. Tao, Z. Huang, Z. Jia, and H. Su, “ManiSkill: Generalizable manipulation skill benchmark with large-scale demonstrations”, in *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2021.
- [12] S. Xu, S. Liu, **T. Mu**, Z. Jia, Y. Wu, and H. Su, “Arena: A scalable and configurable benchmark for policy learning”, *preprint*, 2021.
- [13] **T. Mu\***, J. Gu\*, Z. Jia, H. Tang, and H. Su, “Refactoring policy for compositional generalizability using self-supervised object proposals”, in *Thirty-fourth Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

- [14] F. Liu, Z. Ling, **T. Mu**, and H. Su, “State alignment-based imitation learning”, in *International Conference on Learning Representations (ICLR)*, 2019.
- [15] X. Liu\*, **T. Mu\***, and H. Su, “Transfer value or policy? a value-centric framework towards transferrable continuous reinforcement learning”, in *Deep Reinforcement Learning Workshop at NeurIPS*, 2018.
- [16] Z. Shen, H. Qian, **T. Mu**, and C. Zhang, “Accelerated doubly stochastic gradient algorithm for large-scale empirical risk minimization.”, in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2017.
- [17] Z. Shen, H. Qian, T. Zhou, and **T. Mu**, “Adaptive variance reducing for stochastic gradient descent.”, in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

## INVITED TALKS

---

- **Stanford Vision and Learning Lab** Jul 2024  
*Topic: Modeling the Physical World for Data-Centric Robotics*
- **Microsoft Research Asia** Apr 2023  
*Topic: Modeling the Physical World for Embodied AI: Environments and Algorithms*
- **Pre-training Robot Learning Workshop at CoRL 2022** Dec 2022  
*Topic: On Pre-Training for Visuo-Motor Control: Revisiting a Learning-from-Scratch Baseline*
- **Stanford Vision and Learning Lab** Dec 2021  
*Topic: Generalizable Manipulation Skill Benchmark with Large-Scale Demonstrations*
- **UC Berkeley Robot Learning Lab** Nov 2021  
*Topic: Generalizable Manipulation Skill Benchmark with Large-Scale Demonstrations*
- **Qualcomm AI Lab** Mar 2020  
*Topic: Task-driven Entity Abstraction from Visual Observations*

## PROFESSIONAL SERVICES

---

### Academic Event Organizer

- SAPIEN ManiSkill Challenge 2021(**Lead Organizer**)
- ICLR 2022 Workshop “Generalizable Policy Learning in the Physical World” (**Lead Organizer**)
- CVPR 2022 Tutorial “Building and Working in Environments for Embodied AI”

### Reviewer

- Conference Reviewer: ICLR, NeurIPS, ICML, CoRL, ICRA, IROS, ICCV, CVPR, AAAI, AISTATS
- Journal Reviewer: RA-L, TMLR

## TEACHING

---

- **Co-Instructor** at UC San Diego Spring 2024  
*CSE 276F: Machine Learning for Robotics*
- **Guest Lecturer** at UC San Diego Winter 2023  
*CSE 291-A00: Machine Learning for Robotics*
- **Consultant Volunteer** at UC San Diego Fall 2020  
*CSE 291-J00: Deep Learning Lab (Computer Vision)*
- **Teaching Assistant** at UC San Diego Fall 2018  
*CSE 152: Introduction to Computer Vision*

## AWARDS AND HONORS

---

- ACM-ICPC (International Collegiate Programming Contest) Asia Regional Contest **Gold Medal** 2015
- China Computer Federation Elite Collegiate Award (top 108 in China) 2016
- Award of Excellence for Stars of Tomorrow Internship Program, Microsoft Research Asia 2017