

Bio

I am currently a PhD Candidate at the University of Hong Kong, supervised by Professor Ping Luo. My research focuses on Embodied Intelligence, Multimodal Large Models, Reinforcement Learning, Robot Control, and Autonomous Driving. I have published more than **20 papers** in top AI conferences/SCI Q1 like NeurIPS, ICML, ICLR, CVPR, RSS, IEEE TNNLS and IEEE TITS, with 1 paper as ICML **Oral Presentation** and 3 as **NeurIPS spotlight** papers. Honored with the **Student Best Paper Award** at ICCAS2020 and nominated for **IEEE IV2021's Best Student Paper**. Recipient of the prestigious **Hong Kong PhD Fellowship (HKPFS)** and several other academic awards. Visiting researcher at **NUS-LinS Lab**, **MIT-IBM Watson AI Lab** and **ETH Zurich Computer Vision Lab**. Active reviewer for NeurIPS, ICML, ICLR, CVPR, IJCAI, and AAAI.

My research interest focus on **Embodied AI, Reinforcement Learning, Representation Learning, Robotic Control** and **Autonomous Driving**.

Education

- **The University of Hong Kong**

PhD candidate in Computer Science;

Aug. 2021 – Present

- In year 2021, ranks the 22th in QS World University Rankings.
- I am supervised by **Prof. Ping Luo**, who was named one of the young innovators by the MIT Technology Review "Innovators Under 35 (MIT TR 35)" Asia Pacific, and co-supervised by **Prof. Xiaoou Tang** (IEEE Fellow, director of MMLAB) and study in the **HKUMMLab**.

- **Tsinghua University (C9, Double First Class)**

China

Master in intelligent vehicle Engineering; GPA: 3.78/4.0; Ranking 17/64

Sep. 2018 – June. 2021

- In year 2020, Tsinghua University ranks the 1st in QS Chinese University Ranking and ranks the 15th in QS international University Ranking.
- I obtained the M.Phil Degree under the supervision of **Prof. Bo Cheng** and **Prof. Shengbo Li** at the **Intelligent Driving Laboratory (iDLAB)**.
- Courses with full marks 4.0: "Algorithm Analysis and Design" "Optimal Control" "Applied Stochastic Processes" "Advanced Machine Learning" "Intelligent Transportation Systems Modeling and Simulation" "Statistical Learning Theory and Applications" "Reinforcement Learning and Control" "Vehicle Control Engineering".

- **Harbin institute of technology (HIT, C9, Double First Class)**

China

Bachelor in Mechanical engineering; GPA: 92.01/100; Ranking 1/30

Sep. 2014 – Jun. 2018

- Graduated with the honor of "Provincial Excellent Graduate" of Harbin institute of technology. (Top 10%)
- Graduated with the honor of "Provincial Triple-A Student Pacesetter". (**Top 0.4%**, 10 out of 2400 students)

Honours and Awards

- **Scholarships**

- Hong Kong PhD Fellowship Scheme (**HKPFS**)
- HKU Presidential PhD Scholar Programme (**HKU-PS**)
- National Scholarship for 2014/2015 academic year (**Top 1%**, 2 out of 135 students in HIT)
- National Scholarship for 2015/2016 academic year (**Top 1%**, 2 out of 135 students in HIT)
- National Scholarship for 2016/2017 academic year (**Top 1%**, 2 out of 135 students in HIT)
- First-Class Scholarship for 2014/2015 academic year (**Top 10%**, 13 out of 135 students in HIT)

- **Academic Competitions Awards**

- **Student Best Paper Award** in the 20th International Conference on Control, Automation and Systems (ICCAS) (**Top 1%** (5/500) among accepted papers from 25 countries)
- **Finalists for the Student Best Paper Award** of IV2021 (**3/450**)
- 2st Prize of National College Student Energy Conservation and Emission Reduction Competition (**Top 5%** in China)
- Meritorious Winner Award in Interdisciplinary Contest in Modeling (**Top 13%** worldwide)

- 2nd Second Prize of National College students Ocean Vehicle Design and Production competition (Top 5% in China)
- 3rd of National College Student Mathematics Competition (**Top 10%** in China)
- 1st Prize (provincial) of National College Student Mathematical Modeling Competition (**Top 10%**, in China)

• Honours

- **Outstanding Graduate** of Tsinghua University
- **Outstanding Thesis Award**, Tsinghua University
- Provincial excellent student award (**Top 1%** in HIT)
- Provincial excellent graduates (**Top 5%** in HIT)
- Provincial Triple-A Student Pacesetter (**Top 3%** in HIT)
- Outstanding League Member (**Top 10%** in HIT)

Publications

1. **Yao Mu**, Junting Chen, et al. “RoboCodeX: Multimodal Code Generation for Robotic Behavior Synthesis.” International Conference on Machine Learning (**ICML 2024**)
2. **Yao Mu**, Qinglong Zhang, et al. “EmbodiedGPT: Vision-Language Pre-Training via Embodied Chain of Thought.” Advances in Neural Information Processing Systems 36 (**NeurIPS2023**, **Spotlight**)
3. **Yao Mu**, Shunyu Yao, et al. “EC²: Emergent Communication for Embodied Control.” The IEEE/CVF Conference on Computer Vision and Pattern Recognition 2023 (**CVPR2023**).
4. **Yao Mu**, Zhiqian Lan, et al. “Neural MPC-based Decision-making Framework for Autonomous Driving in Multi-Lane Roundabout.” The 26th IEEE International Conference on Intelligent Transportation Systems(**ITSC 2023**, **Oral Presentation**). This paper was also present in the Scene Representations For Autonomous Driving Workshop on the eleventh International Conference on Learning Representations (**ICLR2023** Autonomous Driving workshop, **Oral Presentation**).
5. **Yao Mu**, Yuzheng Zhuang, et al. “DOMINO: Decomposed Mutual Information Optimization for Generalized Context in Meta-Reinforcement Learning.” Advances in Neural Information Processing Systems 35 (**NeurIPS2022**, **Spotlight**).
6. **Yao Mu**, Shoufa Chen, et al. “CtrlFormer: Learning Transferable State Representation for Visual Control via Transformer.” International Conference on Machine Learning (**ICML 2022**).
7. **Yao Mu**, Yuzheng Zhuang, et al. “Model-Based Reinforcement Learning via Imagination with Derived Memory.” Advances in Neural Information Processing Systems 34 (**NeurIPS2021**).
8. **Yao Mu**, Baiyu Peng, et al.. “Mixed reinforcement learning for efficient policy optimization in stochastic environments.” 2020 20th International Conference on Control, Automation and Systems (ICCAS). IEEE, 2020. (**Student Best Paper Award**) .
9. Pengying Wu*, **Yao Mu***(**Co-first author**), et al. “VoroNav: Voronoi-based Zero-shot Object Navigation with Large Language Model” International Conference on Machine Learning (**ICML 2024**).
10. Zeyu Gao*, **Yao Mu***(**Co-first author**), et al. “SEM2: Enhance Sample Efficiency and Robustness of End-to-end Urban Autonomous Driving via Semantic Masked World Model.” IEEE Transactions on Intelligent Transportation Systems (**IEEE TITS**, **Impact Factor:9.551**).
11. Shentao Qin, **Yao Mu**, et al. “Feasible Reachable Policy Iteration.” International Conference on Machine Learning (**ICML 2024**).
12. Yuwei Zeng, **Yao Mu**, et al. “Learning Reward for Robot Skills Using Large Language Models via Self-Alignment.” International Conference on Machine Learning (**ICML 2024**).
13. Zhixuan Liang, **Yao Mu**, et al. “AdaptDiffuser: Diffusion Models as Adaptive Self-evolving Planners.” The IEEE/CVF Conference on Computer Vision and Pattern Recognition 2024 (**CVPR2024**).
14. Zhixuan Liang, **Yao Mu**, et al. “Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution.” International Conference on Machine Learning (**ICML 2023**, **Oral Presentation**).
15. Mengkang Hu, **Yao Mu**, et al. “Tree-planner: Efficient close-loop task planning with large language models”. The International Conference on Learning Representations (**ICLR2024**).

16. Fei Ni, Jianye Hao, **Yao Mu**, et al. “MetaDiffuser: Diffusion Model as Conditional Planner for Offline Meta-RL.” International Conference on Machine Learning (**ICML 2023**).
17. Runjian Chen, **Yao Mu**, et al. “CO3: Cooperative Unsupervised 3D Representation Learning for Autonomous Driving”. The International Conference on Learning Representations (**ICLR2023**).
18. Zibin Dong, Yifu Yuan, Jianye Hao, Fei Ni, **Yao Mu**, et al. “AlignDiff: Aligning Diverse Human Preferences via Behavior-Customisable Diffusion Model”. The International Conference on Learning Representations (**ICLR2024**).
19. Yifu Yuan, Jianye Hao, Fei Ni, **Yao Mu**, et al. “EUCLID: Towards Efficient Unsupervised Reinforcement Learning with Multi-choice Dynamics Model”. The International Conference on Learning Representations (**ICLR2023**).
20. Zeyu Gao*, **Yao Mu***(**co-first author**), et al. “SEM2: Enhance Sample Efficiency and Robustness of End-to-end Urban Autonomous Driving via Semantic Masked World Model.” Advances in Neural Information Processing Systems 35 (**NeurIPS2022** Deep RL workshop, Under review at IEEE TITS).
21. Yao Lai, **Yao Mu**, et al. “MaskPlace: Fast Chip Placement via Reinforced Visual Representation Learning.” Advances in Neural Information Processing Systems 35 (**NeurIPS2022**, **Spotlight**).
22. Xiaoyu Chen, **Yao Mu**, et al. “Flow-based Recurrent Belief State Learning for POMDPs.” International Conference on Machine Learning (**ICML 2022**).
23. Qiushan Guo, **Yao Mu**, et al. “Scale-Equivalent Distillation for Semi-Supervised Object Detection.” Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.(**CVPR2022**).
24. Zhecheng Yuan, Guozheng Ma, **Yao Mu**, et al. “Don’t Touch What Matters: Task-Aware Lipschitz Data Augmentation for Visual Reinforcement Learning.” International Joint Conferences on Artificial Intelligence (**IJCAI2022**): 3702-3708
25. Mingxiao Huo, et al., **Yao Mu**, et al. “Human-oriented Representation Learning for Robotic Manipulation.” The 20th conference of the “Robotics: Science and Systems” (**RSS2024**)
26. Baiyu Peng, **Yao Mu**, et al. “Model-based actor-critic with chance constraint for stochastic system.” 2021 60th IEEE Conference on Decision and Control (**CDC**). IEEE, 2021.
27. Baiyu Peng, **Yao Mu**, et al. “Separated proportional-integral lagrangian for chance constrained reinforcement learning.” 2021 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2021. (**Finalist of the Best Student Paper Award**)
28. Baiyu Peng, **Yao Mu**, et al. “Model-based Chance-Constrained Reinforcement Learning via Separated Proportional-Integral Lagrangian.” IEEE Transactions on Neural Networks and Learning Systems(**IEEE TNNLS**, **Impact Factor: 10.451**).
29. Junjie Wang, Qichao Zhang, **Yao Mu**, et al. “Prototypical context-aware dynamics generalization for high-dimensional model-based reinforcement learning.” IEEE Transactions on Industrial Informatics(**IEEE TII**, **Impact Factor: 10.215**).
30. Yuhang Zhang, **Yao Mu**, et al. “Steadily Learn to Drive with Virtual Memory.” 2022 11th Asia-Pacific Regional Conference of the ISTVS (ISTVS 2022).
31. Dafeng Chi, Yuzheng zhuang, **Yao Mu**, et al. “Offline-to-online Co-evolutional User Simulator and DIALOGUE System.” The Sere TOD workshop on the 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022** Sere TOD workshop).

Skills

- **Programming:** Python(Expert), Pytorch(Expert), C++(Intermediate), TensorFlow(Intermediate)
- **Softwares:** Linux, Git, ROS, Pycharm, Visual Studio, Microsoft Office
- **Github:** <https://github.com/YaoMarkMu>

Contact

- **Email:** muyao@connect.hku.hk
- **Personal homepage:** <https://yaomarkmu.github.io>
- **Twitter:** <https://mobile.twitter.com/YaoMarkMu1>