Zhenghao PENG

Homepage: https://pengzhenghao.github.io/

Google Scholar
Github: pengzhenghao
Email: pzh@cs.ucla.edu

RESEARCH INTERESTS

Robotics, Vision-Language-Action (VLA) Model, Human-in-the-loop Learning, (Multi-agent) Simulation, and Large-scale Reinforcement Learning.

My research focuses on building physically interactive agents that can reason, align with human intent, and generalize across diverse tasks. I aim at building foundation models for robotics—scalable, aligned, and deployable in the real world to improve human well-being.

A detailed **Research Statement** is available on my homepage.

Robots Experience: Unitree Go2 Quadruped Robot, Coco Delivery Robot (Four-Wheeled), Marine Robots.

EDUCATION

University of California, Los Angeles (UCLA) *September 2022 - May 2026 (Expected)*

- PhD student at the Department of Computer Science, supervised by Prof. Bolei Zhou.

The Chinese University of Hong Kong (CUHK) *August 2019 - July 2022*

- MPhil student at the Department of Information Engineering, supervised by Prof. Bolei Zhou.

Shanghai Jiao Tong University (SJTU) *Sept. 2015 - July 2019*

- Bachelor of Engineering and member of Zhiyuan Honors Program.
- Research assistant supervised by Prof. Li Jiang.

EXPERIENCE

NVIDIA, Santa Clara, CA *June 2025 - Present*

- Research intern in Autonomous Vehicle Group at NVIDIA Research. Mentored by Boris Ivanovic.
- Developing VLA models for autonomous driving, focusing on embodied reasoning and action alignment.

Waymo LLC, Mountain View, CA *June 2023 - September 2023*

- Research intern in behavior modeling. Mentored by Justin Fu and Rowan McAllister.
- Improved the average displacement error of multi-agent behavior models by 25% using RL as a closed-loop post-training technique.

Shenzhen Institutes of Advanced Technology (SIAT), Chinese Academy of Sciences *June 2018 - Sept. 2018*

- Research intern at Multimedia Research Center, under the supervision of Prof. Yu Qiao.

OPEN-SOURCE PROJECTS

- I am a strong advocate of open-source and reproducible research. Most of my works are released with public code, and my contributions can be found on GitHub.
- Notably, MetaDrive is an open-source driving simulator for reinforcement learning and autonomous driving. It has received ∼1000 GitHub stars and 350+ citations, and is widely adopted in research community.

AWARDS AND HONORS

- Dissertation Year Award 2025-2026, UCLA
- Amazon Fellowship, 2024-2025, UCLA
- University Fellowship, 2023-2024, 2024-2025, UCLA
- The Outstanding Tutors Award 2021 of the Faculty of Engineering , 2021, CUHK
- Teaching Assistant Award, Term 1 2020 - 2021, Term 2 2020 - 2021, CUHK

- Postgraduate Studentship, 2019 - 2022, CUHK
- Zhiyuan Honors Scholarship, 2015 - 2017, SJTU

RESEARCH PAPERS

[1] **Zhenghao Peng**, Yuxin Liu, and Bolei Zhou. Infgen: Scenario generation as next-token-group prediction. In *Preprint*, 2025 [Webpage, PDF]

[2] Yuxin Liu*, **Zhenghao Peng***, Xuanhao Cui, and Bolei Zhou. Adv-bmt: Bidirectional motion transformer for safety-critical traffic scenario generation. In *Preprint*, 2025 [Webpage, PDF]

[3] **Zhenghao Peng**, Zhizheng Liu, and Bolei Zhou. Data-efficient learning from human interventions for mobile robots. *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025 (**ICRA 2025**) [Webpage, PDF]

[4] **Zhenghao Peng**, Wenjie Luo, Yiren Lu, Tianyi Shen, Cole Gulino, Ari Seff, and Justin Fu. Improving agent behaviors with rl fine-tuning for autonomous driving. *European Conference on Computer Vision*, 2024 (**ECCV 2024**)[PDF]

[5] Yunsong Zhou, Michael Simon, **Zhenghao Peng**, Sicheng Mo, Hongzi Zhu, Minyi Guo, and Bolei Zhou. Simgen: Simulator-conditioned driving scene generation. *Advances in Neural Information Processing Systems*, 2024 (**NeurIPS 2024**) [PDF, Webpage]

[6] Brandon J. McMahan, **Zhenghao Peng**, Bolei Zhou, and Jonathan C. Kao. Shared autonomy with ida: Interventional diffusion assistance. *Advances in Neural Information Processing Systems*, 2024 (**NeurIPS 2024**) [PDF]

[7] **Zhenghao Peng**, Wenjie Mo, Chenda Duan, Quanyi Li, and Bolei Zhou. Learning from active human involvement through proxy value propagation. *Advances in Neural Information Processing Systems*, 2023 (**NeurIPS 2023 *Spotlight***) [PDF, Webpage]

[8] Quanyi Li*, **Zhenghao Peng***, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and Bolei Zhou. Scenarionet: Open-source platform for large-scale traffic scenario simulation and modeling. *Advances in Neural Information Processing Systems*, 2023 (**NeurIPS 2023**) [PDF , Code , Webpage]

[9] Linrui Zhang, **Zhenghao Peng**, Quanyi Li, and Bolei Zhou. Cat: Closed-loop adversarial training for safe end-to-end driving. In *7th Annual Conference on Robot Learning*, 2023 (**CoRL 2023**) [PDF, Code, Webpage]

[10] Lan Feng*, Quanyi Li*, **Zhenghao Peng***, Shuhan Tan, and Bolei Zhou. Trafficgen: Learning to generate diverse and realistic traffic scenarios. In *2023 International Conference on Robotics and Automation (ICRA)*. IEEE, 2023 (**ICRA 2023**) [PDF, Code, Webpage]

[11] Zhenghai Xue, **Zhenghao Peng**, Quanyi Li, Zhihan Liu, and Bolei Zhou. Guarded policy optimization with imperfect online demonstrations. In *International Conference on Learning Representations*, 2023 (**ICLR 2023**) [PDF, Code, Webpage]

[12] Quanyi Li, **Zhenghao Peng**, Haibin Wu, Lan Feng, and Bolei Zhou. Human-AI shared control via policy dissection. *Advances in Neural Information Processing Systems*, 2022 (**NeurIPS 2022**)[PDF, Code, Webpage]

[13] Qihang Zhang, **Zhenghao Peng**, and Bolei Zhou. Learning to drive by watching youtube videos: Action-conditioned contrastive policy pretraining. *European Conference on Computer Vision*, 2022 (**ECCV 2022**)[PDF, Webpage]

[14] Quanyi Li*, **Zhenghao Peng***, Zhenghai Xue, Qihang Zhang, and Bolei Zhou. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE transaction on Pattern Analysis and Machine Intelligence*, 2021 (**TPAMI**) [Paper, Code, Webpage]

[15] Mingxin Huang, Yuliang Liu, **Zhenghao Peng**, Chongyu Liu, Dahua Lin, Shenggao Zhu, Nicholas Yuan, Kai Ding, and Lianwen Jin. Swintextspotter: Scene text spotting via better synergy between text detection and text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022 (**CVPR 2022**)

[16] Quanyi Li*, **Zhenghao Peng***, and Bolei Zhou. Efficient learning of safe driving policy via human-AI copilot optimization. In *International Conference on Learning Representations*, 2022 (**ICLR 2022**) [PDF, Code, Webpage]

[17] **Zhenghao Peng***, Quanyi Li*, Chunxiao Liu, and Bolei Zhou. Safe driving via expert guided policy optimization. In *5th Annual Conference on Robot Learning*, 2021 (**CoRL 2021**) [PDF, Code, Webpage, Poster]

[18] **Zhenghao Peng**, Quanyi Li, Ka Ming Hui, Chunxiao Liu, and Bolei Zhou. Learning to simulate self-driven particles system with coordinated policy optimization. *Advances in Neural Information Processing Systems*, 34, 2021 (**NeurIPS 2021**) [PDF, Code, Webpage, Poster]

[19] Quanyi Li*, **Zhenghao Peng***, Qihang Zhang, Chunxiao Liu, and Bolei Zhou. Improving the generalization of end-to-end driving through procedural generation. *arXiv preprint arXiv:2012.13681*, 2020 [PDF, Repo, Webpage]

[20] **Zhenghao Peng**, Hao Sun, and Bolei Zhou. Non-local policy optimization via diversity-regularized collaborative exploration. *arXiv preprint arXiv:2006.07781*, 2020 [PDF]

[21] Zhuoran Song, Dongyu Ru, Ru Wang, Hongru Huang, **Zhenghao Peng**, Jing Ke, Xiaoyao Liang, and Li Jiang. Approximate random dropout. In *Design, Automation & Test in Europe Conference & Exhibition, 2019. DATE'19*. IEEE, 2019 [PDF]

[22] **Zhenghao Peng**, Xuyang Chen, Chengwen Xu, Naifeng Jing, Xiaoyao Liang, Cewu Lu, and Li Jiang. Axnet: Approximate computing using an end-to-end trainable neural network. In *Proceedings of the 2018 International Conference on Computer-Aided Design. ICCAD'18*. IEEE/ACM, 2018 [PDF]

TALKS

- Human-in-the-loop Agent Learning, EECS 598: Action and Perception Guest Lecture, invited by: Stella Yu, May 2024

TEACHING EXPERIENCES

- CS260R Reinforcement Learning, UCLA, 2025 Winter
- CS260R Reinforcement Learning, UCLA, 2023 Fall
- CS269 Seminar on Reinforcement Learning, UCLA, 2022 Fall
- IERG5350 Reinforcement Learning, CUHK, Term 1, 2021-22
- CSCI2100E Data Structures, CUHK, Term 2, 2020-21
- IERG5350 Reinforcement Learning, CUHK, Term 1, 2020-21
- IERG6130 Seminar on Reinforcement Learning, CUHK, Term 2, 2019-20

MISCELLANEOUS

- **Reviewer:** NeurIPS, ICML, CVPR, ECCV, ICLR, CoRL, RSS, IROS, ICRA, AAAI, TNNLS, IJCV, ICCV, RA-L, *etc.*
- **Frameworks:** PyTorch, Jax, TensorFlow, Ray, RLLib, Keras, ROS2, *etc.*
- **Skills:** LaTeX, Keynote, Photoshop, Final Cut Pro, Git, Cantonese, Photography, *etc.*