

Shiyu Huang (黄世宇)

Haidian District, Beijing, China, 100084

huangsy1314@163.com

[\[OpenRL\]](#) [\[知乎\]](#) [\[GitHub\]](#) [\[Personal Website\]](#) [\[Google Scholar\]](#) [\[Linkedin\]](#)

updated in 6/2026

Biography

I am a principal researcher at XPeng Inc. Before that, I worked as a principal researcher in Zhipu AI and 4Paradigm Inc., and I am the leader of the OpenRL Lab. I received my B.E. and Ph.D. degrees (co-advised by Prof. Jun Zhu and Prof. Ting Chen) from the Department of Computer Science and Technology, Tsinghua University in July 2017 and June 2022. My researches focus on large language models, multimodal large language models, agents, and deep reinforcement learning. I have also worked at 4Paradigm Inc., RealAI Inc., Huawei Noah's Ark Lab, Tencent AI Lab, Carnegie Mellon University, and SenseTime Inc.. And I am also the founder of the OpenRL Lab (1.8k+ GitHub stars).

Education

- **Tsinghua University**(清华大学), Beijing, China
Doctor of Philosophy (PhD)(博士), Computer Science and Technology, co-advised by Prof. **Jun Zhu** and Prof. **Ting Chen** Sept. 2017 - Jun. 2022
- **Tsinghua University**(清华大学), Beijing, China
Bachelor of Engineering(本科), Computer Science and Technology Aug. 2013 - Jul. 2017
 - cumulative GPA: **90/100**
 - Tsinghua Excellent Graduates (**top10%**)
 - won Academic Excellence Award every year
 - 四川省 2013 年高考理科前 30 名

Work Experience

- **Technical Manager**, XPENG(小鹏汽车), Beijing, China Jun. 2025 - Now
 - Team Leader, Foundation Model
 - LLM: Pre-train, SFT, MTP
 - Agents: General Agent, GUI Agent
 - Audio LM: ASR, TTS
 - Visual LM: Visual Grounding, Spatial Understanding, OCR
 - Omni-model: unify text, audio, and visual modalities.
- **Researcher**, Zhipu AI(智谱 AI), Beijing, China Mar. 2024 - Jun. 2025
 - Team Leader, Project Manager
 - **Video Call**: Advanced real-time video call model.
 - **GLM-4V-Plus**: A commercial-grade general video understanding model.

- **CogVLM2-Video**: A Temporally-Aware Video Understanding Model
- **CogVideoX**: Text-to-Video Diffusion Models with An Expert Transformer
- **LVBench**: An Extreme Long Video Understanding Benchmark
- **Research Scientist**, 4Paradigm(第四范式), Beijing, China Jul. 2022 - Mar. 2024
 - Founder and leader of OpenRL Lab: <https://github.com/OpenRL-Lab/>.
 - Leader of developing distributed multi-agent reinforcement learning framework.
 - Leader of developing self-play reinforcement learning framework.
 - TiZero: Mastering Multi-Agent Football with Curriculum Learning and Self-Play.
- **Researcher**, RealAI(瑞莱智慧), Beijing, China Jun. 2021 - Apr. 2022
 - Leader of developing distributed multi-agent reinforcement learning framework.
 - Leader of developing offline reinforcement learning benchmarks.
 - TiKick: Towards Playing Multi-agent Football Full Games from Single-agent Demonstrations.
- **Research Intern**, Huawei Noah’s Ark Lab(华为诺亚), Beijing, China May. 2020 - Jun. 2021
 - Develop algorithms for circuit routing. The work is submitted to AAMAS 2022.
 - Do research for Deep Reinforcement Learning algorithms. The work is published on PRICAI 2021 and CoG 2022.
- **Research Intern**, Tencent(腾讯) AI Lab, Shen Zhen, China Apr. 2018 - Feb. 2019
 - Developed an AI agent for ViZDoom. The work is published on AAAI 2019.
- **Research Intern**, Carnegie Mellon University(卡内基梅隆大学), Pittsburgh, USA Jun. 2016 - Sep. 2016
 - Developed a detection framework. The work is published on CVPR 2017.
- **Research Intern**, SenseTime(商汤), Beijing, China Jun. 2015 - Dec. 2015
 - Vehicle Recognition and Face Verification.

Publications & Preprints

(* equal contribution)

- **OmniGUI: Benchmarking GUI Agents in Omni-Modal Smartphone Environments**
Felix Henry, Xiaochen Lin, Jiangyou Zhu, Yangfan, Bingqian Zhang, Min Chen, and **Shiyu Huang**
arXiv:2605.18758, 2026
- **ME-IGM: Individual-Global-Max in Maximum Entropy Multi-Agent Reinforcement Learning**
Wentse Chen, Yuxuan Li, **Shiyu Huang**, Jiayu Chen, Jeff Schneider
The 25th International Conference on Autonomous Agents and Multi-Agent Systems(**AAMAS**), 2026

- [OmniDPO: A Preference Optimization Framework to Address Omni-Modal Hallucination](#)
Junzhe Chen, Tianshu Zhang, **Shiyu Huang**, Yuwei Niu, Chao Sun, Rongzhou Zhang, Guanyu Zhou, Lijie Wen, Xuming Hu
The Fortieth AAAI Conference on Artificial Intelligence(AAAI), 2026
- [Can Large Language Models Master Complex Card Games?](#)
Wei Wang, Fuqing Bie, Junzhe Chen, Dan Zhang, **Shiyu Huang**, Evgeny Kharlamov, Jie Tang
The Thirty-Ninth Annual Conference on Neural Information Processing Systems(NeurIPS), 2025
- [A Survey on Parallel Text Generation: From Parallel Decoding to Diffusion Language Models](#)
Lingzhe Zhang*, Liancheng Fang*, Chiming Duan*, Minghua He*, Leyi Pan*, Pei Xiao, **Shiyu Huang**, Yunpeng Zhai, Xuming Hu, Philip S. Yu, Aiwei Liu
arXiv:2508.08712, 2025
- [Omni-SafetyBench: A Benchmark for Safety Evaluation of Audio-Visual Large Language Models](#)
Leyi Pan, Zheyu Fu, Yunpeng Zhai, Shuchang Tao, Sheng Guan, **Shiyu Huang**, Lingzhe Zhang, Zhaoyang Liu, Bolin Ding, Felix Henry, Lijie Wen, Aiwei Liu
arXiv:2508.07173, 2025
- [OmniPlay: Benchmarking Omni-Modal Models on Omni-Modal Game Playing](#)
Fuqing Bie, **Shiyu Huang**, Xijia Tao, Zhiqin Fang, Leyi Pan, Junzhe Chen, Min Ren, Liuyu Xiang, Zhaofeng He
arXiv:2508.04361, 2025
- [SA-GCS: Semantic-Aware Gaussian Curriculum Scheduling for UAV Vision-Language Navigation](#)
Hengxing Cai*, Jinhan Dong*, Yijie Rao*, Jingcheng Deng, Jingjun Tan, Qien Chen, Haidong Wang, Zhen Wang, **Shiyu Huang**, Agachai Sumalee, Renxin Zhong
arXiv:2508.00390, 2025
- [GLM-4.1V-Thinking: Towards Versatile Multimodal Reasoning with Scalable Reinforcement Learning](#)
Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihang Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, Aohan Zeng, Baoxu Wang, Boyan Shi, Changyu Pang, Chenhui Zhang, Da Yin, Fan Yang, Guoqing Chen, Jiazheng Xu, Jiali Chen, Jing Chen, Jinhao Chen, Jinghao Lin, Jinjiang Wang, Junjie Chen, Leqi Lei, Letian Gong, Leyi Pan, Mingzhi Zhang, Qinkai Zheng, Sheng Yang, Shi Zhong, **Shiyu Huang**, Shuyuan Zhao, Siyan Xue, Shangqin Tu, Shengbiao Meng, Tianshu Zhang, Tianwei Luo, Tianxiang Hao, Wenkai Li, Wei Jia, Xin Lyu, Xuancheng Huang, Yanling Wang, Yadong Xue, Yanfeng Wang, Yifan An, Yifan Du, Yiming Shi, Yiheng Huang, Yilin Niu, Yuan Wang, Yuanchang Yue, Yuchen Li, Yutao Zhang, Yuxuan Zhang, Zhanxiao Du, Zhenyu Hou, Zhao Xue, Zhengxiao Du, Zihan Wang, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Minlie Huang, Yuxiao Dong, Jie Tang
arXiv:2507.01006, 2025
- [Generalizable Agent Modeling for Agent Collaboration-Competition Adaptation with Multi-Retrieval and Dynamic Generation](#)
Chenxu Wang, Yonggang Jin, Cheng Hu, Youpeng Zhao, Zipeng Dai, Jian Zhao, **Shiyu Huang**, Liuyu Xiang, Junge Zhang, Zhaofeng He
Neurocomputing (2025): 130912
- [MotionBench: Benchmarking and Improving Fine-grained Video Motion Understanding for Vision Language Models](#)
Wenyi Hong*, Yean Cheng*, Zhuoyi Yang, Weihang Wang, Lefan Wang, Xiaotao Gu, **Shiyu Huang**, Yuxiao Dong, Jie Tang

The IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR), 2025

- [Can LLM Watermarks Robustly Prevent Unauthorized Knowledge Distillation?](#)
Leyi Pan, Aiwei Liu, **Shiyu Huang**, Yijian Lu, Xuming Hu, Lijie Wen, Irwin King, Philip S. Yu
The 63rd Annual Meeting of the Association for Computational Linguistics (ACL), 2025
- [VisionReward: Fine-Grained Multi-Dimensional Human Preference Learning for Image and Video Generation](#)
Jiazheng Xu*, Yu Huang*, Jiale Cheng, Yuanming Yang, Jiajun Xu, Yuan Wang, Wenbo Duan, Shen Yang, Qunlin Jin, Shurun Li, Jiayan Teng, Zhuoyi Yang, Wendi Zheng, Xiao Liu, Ming Ding, Xiaohan Zhang, Xiaotao Gu, **Shiyu Huang**, Minlie Huang, Jie Tang, Yuxiao Dong
The Fortieth AAAI Conference on Artificial Intelligence(AAAI), 2026
- [ICT: Image-Object Cross-Level Trusted Intervention for Mitigating Object Hallucination in Large Vision-Language Models](#)
Junzhe Chen*, Tianshu Zhang*, **Shiyu Huang**, Yuwei Niu, Linfeng Zhang, Lijie Wen, Xuming Hu
The IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR), 2025
- [DreamPolish: Domain Score Distillation With Progressive Geometry Generation](#)
Yean Cheng, Ziqi Cai, Ming Ding, Wendi Zheng, **Shiyu Huang**, Yuxiao Dong, Jie Tang, Boxin Shi
arXiv:2411.01602, 2024
- [CogVLM2: Visual Language Models for Image and Video Understanding](#)
Wenyi Hong, Weihang Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng, **Shiyu Huang**, Junhui Ji, Zhao Xue, Lei Zhao, Zhuoyi Yang, Xiaotao Gu, Xiaohan Zhang, Guanyu Feng, Da Yin, Zihan Wang, Ji Qi, Xixuan Song, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Yuxiao Dong, Jie Tang
arXiv:2408.16500, 2024
- [CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer](#)
Zhuoyi Yang*, Jiayan Teng*, Wendi Zheng, Ming Ding, **Shiyu Huang**, Jiazheng Xu, Yuanming Yang, Xiaohan Zhang, Xiaotao Gu, Guanyu Feng, Da Yin, Wenyi Hong, Weihang Wang, Yean Cheng, Yuxuan Zhang, Ting Liu, Bin Xu, Yuxiao Dong, Jie Tang
The Thirteenth International Conference on Learning Representations (ICLR), 2025
- [WaterSeeker: Pioneering Efficient Detection of Watermarked Segments in Large Documents](#)
Leyi Pan, Aiwei Liu, Yijian Lu, Zitian Gao, Yichen Di, **Shiyu Huang**, Lijie Wen, Irwin King, Philip S. Yu
The AAAI 2025 PDLM Workshop on Preventing and Detecting LLM Misinformation, 2025
- [A Survey on Self-play Methods in Reinforcement Learning](#)
Ruize Zhang, Zelai Xu, Chengdong Ma, Chao Yu, Wei-Wei Tu, **Shiyu Huang**, Deheng Ye, Wenbo Ding, Yaodong Yang, Yu Wang
arXiv:2408.01072, 2024
- [Soft-QMIX: Integrating Maximum Entropy For Monotonic Value Function Factorization](#)
Wentse Chen, **Shiyu Huang**, Jeff Schneider
arXiv:2406.13930, 2024
- [LVBench: An Extreme Long Video Understanding Benchmark](#)
Weihang Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, **Shiyu Huang**, Bin Xu, Yuxiao Dong, Ming Ding, Jie Tang
International Conference on Computer Vision(ICCV)(**Highlight**), 2025

- [MQE: Unleashing the Power of Interaction with Multi-agent Quadruped Environment](#)
Ziyang Xiong, Bo Chen, **Shiyu Huang**, Wei-Wei Tu, Zhaofeng He, Yang Gao
The 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (**IROS**), 2024
- [LLMArena: Assessing Capabilities of Large Language Models in Dynamic Multi-Agent Environments](#)
Junzhe Chen, Xuming Hu, Shuodi Liu, **Shiyu Huang**, Wei-Wei Tu, Zhaofeng He, Lijie Wen
The 62nd Annual Meeting of the Association for Computational Linguistics(**ACL**), 2024
- [AutoSAT: Automatically Optimize SAT Solvers via Large Language Models](#)
Yiwen Sun, Xianyin Zhang, **Shiyu Huang**, Shaowei Cai, Bing-Zhen Zhang, Ke Wei
The Journal of Artificial Intelligence Research, 2026
- [OpenRL: A Unified Reinforcement Learning Framework](#)
Shiyu Huang, Wentse Chen, Yiwen Sun, Fuqing Bie, Wei-Wei Tu
arXiv:2312.16189, 2023
- [DGPO: Discovering Multiple Strategies with Diversity-Guided Policy Optimization](#)
Wenze Chen, **Shiyu Huang**, Yuan Chiang, Tim Pearce, Wei-Wei Tu, Ting Chen, Jun Zhu
Thirty-Eighth AAAI Conference on Artificial Intelligence(**AAAI**), Vancouver, Canada, 2024
- [SwiftSage: A Generative Agent with Fast and Slow Thinking for Complex Interactive Tasks](#)
Bill Yuchen Lin, Yicheng Fu, Karina Yang, Prithviraj Ammanabrolu, Faeze Brahman, **Shiyu Huang**,
Chandra Bhagavatula, Yejin Choi, Xiang Ren
Thirty-seventh Conference on Neural Information Processing Systems(**NeurIPS**)(**Spotlight**), 2023
- [Robustness and Generalizability of Deepfake Detection: A Study with Diffusion Models](#)
Haixu Song, **Shiyu Huang**, Yinpeng Dong, Wei-Wei Tu
arXiv:2309.02218, 2023
- [Diverse Policies Converge in Reward-free Markov Decision Processes](#)
Fanqi Lin, **Shiyu Huang**, Wei-Wei Tu
The 20th Pacific Rim International Conference on Artificial Intelligence(**PRICAI**), Jakarta, Indonesia, 2023
- [Uncertainty quantification via a memristor Bayesian deep neural network for risk-sensitive reinforcement learning](#)
Yudeng Lin, Qingtian Zhang, Bin Gao, Jianshi Tang, Peng Yao, Chongxuan Li, **Shiyu Huang**,
Zhengwu Liu, Ying Zhou, Yuyi Liu, Wenqiang Zhang, Jun Zhu and He Qian
Nature Machine Intelligence, 2023
- [TiZero: Mastering Multi-Agent Football with Curriculum Learning and Self-Play](#)
Fanqi Lin*, **Shiyu Huang***, Tim Pearce, Wenze Chen and Wei-Wei Tu
The 22nd International Conference on Autonomous Agents and Multiagent Systems(**AAMAS**),
London, UK, 2023
- [Learning Graph-Enhanced Commander-Executor for Multi-Agent Navigation](#)
Xinyi Yang, **Shiyu Huang**, Yiwen Sun, Yuxiang Yang, Chao Yu, Wei-Wei Tu, Huazhong Yang and Yu
Wang
The 22nd International Conference on Autonomous Agents and Multiagent Systems(**AAMAS**),
London, UK, 2023
- [DGPO: Discovering Multiple Strategies with Diversity-Guided Policy Optimization](#)
Wenze Chen, **Shiyu Huang**, Yuan Chiang, Ting Chen, Jun Zhu
The 22nd International Conference on Autonomous Agents and Multiagent Systems(**AAMAS**)
Extended Abstract, London, UK, 2023
- [VMAPD: Generate Diverse Solutions for Multi-Agent Games with Recurrent Trajectory Discriminators](#)

Shiyu Huang*, Chao Yu*, Bin Wang, Dong Li, Yu Wang, Ting Chen and Jun Zhu
IEEE Conference on Games (**CoG**), Beijing, China, 2022. (**Best Paper Nomination**)

- **TiKick: Towards Playing Multi-agent Football Full Games from Single-agent Demonstrations**
Shiyu Huang*, Wenze Chen*, Longfei Zhang, Shizhen Xu, Ziyang Li, Fengming Zhu, Deheng Ye, Ting Chen and Jun Zhu
NeurIPS-21 Workshop: 2nd Offline Reinforcement Learning Workshop.
- **Ranking Cost: Building An Efficient and Scalable Circuit Routing Planner with Evolution-Based Optimization**
Shiyu Huang, Bin Wang, Dong Li, Jianye Hao, Jun Zhu and Ting Chen
IJCAI-ECAI 2022 Workshop: The 2nd International Workshop on Heuristic Search in Industry, Vienna, Austria, 2022.
- **Deep Reinforcement Learning with Credit Assignment for Combinatorial Optimization**
Dong Yan, Jiayi Weng, **Shiyu Huang**, Chongxuan Li, Yichi Zhou, Hang Su, Jun Zhu
Pattern Recognition, 2022.
- **Off-policy Training for Truncated TD(λ) Boosted Soft Actor-Critic**
Shiyu Huang, Bin Wang, Hang Su, Dong Li, Jianye Hao, Jun Zhu and Ting Chen
The Pacific Rim International Conference on Artificial Intelligence (**PRICAI**), Hanoi, Vietnam, 2021.
- **SVQN: Sequential Variational Soft Q-Learning Networks**
Shiyu Huang, Hang Su, Jun Zhu, and Ting Chen
Eighth International Conference on Learning Representations (**ICLR**), Millennium Hall, Addis Ababa ETHIOPIA, 2020
- **Combo-Action Training Agent For FPS Game with Auxiliary Tasks**
Shiyu Huang, Hang Su, Jun Zhu, and Ting Chen
The Thirty-Third AAAI Conference on Artificial Intelligence (**AAAI**)(**Spotlight**), Honolulu, Hawaii, USA, 2019
- **Expecting the Unexpected: Training Detectors for Unusual Pedestrians with Adversarial Imposters**
Shiyu Huang, and Deva Ramanan
IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), Honolulu, Hawaii, USA, 2017
- **Learning to Assign Credit in Reinforcement Learning by Incorporating Abstract Relations**
Dong Yan, **Shiyu Huang**, Hang Su, and Jun Zhu
AAAI-19 Workshop on Reinforcement Learning in Games
- **Model-based Credit Assignment for Model-free Deep Reinforcement Learning**
Dong Yan, Jiayi Weng, **Shiyu Huang**, Chongxuan Li, Yichi Zhou, Hang Su, Jun Zhu

Talks

- **RLHF @ Zhiyuan Community**, 2023.8.
 - slide: <https://drive.google.com/file/d/1qh92DrnsN5yjEbk43nG4hlw2cZd1hnlub/view>
- **OpenRL @ 5th BAAI Conference**, 2023.6.
 - video: https://www.youtube.com/watch?v=3_EtNuLJRdM

Projects

- **CogVideoX**.

- 12.8k stars, 1.3k forks
- <https://github.com/THUDM/CogVideo>
- **CogVLM2.**
 - 2.4k stars, 163 forks
 - <https://github.com/THUDM/CogVLM2>
- **OpenRL: Unified reinforcement learning framework.**
 - 835 stars, 81 forks
 - <https://github.com/OpenRL-Lab/openrl>
- **TiZero: RL agents for Google Research Football.**
 - <https://github.com/OpenRL-Lab/TiZero>
- **ChatAgent: A Python-based agent framework for large language models.**
 - <https://github.com/OpenRL-Lab/ChatAgent>
- **TLaunch: Launch programs on multiple hosts.**
 - <https://github.com/TARTRL/TLaunch>
- **OpenPlugin: Toolkit to manage the plugins of the large language model.**
 - <https://github.com/OpenRL-Lab/OpenPlugin>

Patents

- **Generation method, device, medium and computing device of diversity strategy.**
Shiyu Huang, Tian Tian. 2021116684627
- **Method or equipment for controlling agent.**
Jun Zhu, Shiyu Huang, Hang Su. ZL201910078546.1

Honors & Awards

- **Tung OOCL Scholarship, Tsinghua University, 2019**
- **Tsinghua Excellent Graduates, Tsinghua University, 2017**
- **Academic Excellence Award, Tsinghua University, 2014-2016**

Competitions

- 2022.8
IEEE CoG 2022 Football AI Competition:
[Track2](#), 3rd place
- 2018.8
ViZDoom 2018 AI Competition:
[Track1](#), 1st place
[Track2](#), 2nd place

- 2017
ViZDoom 2017 AI Competition:
[Track2](#), 2nd place

Skills

- **Programming:** C++/C, Python, Matlab, Java, Ruby, C#
- **OS:** Linux, MacOS, Windows
- **Software:** Tensorflow, Pytorch, OpenCV, Flask, Unity3D, Qt5.

Services

Organizer for:

NeurIPS 2023 Workshop on New in ML:

<https://nehzux.github.io/NewInML2023NeurIPS/newinml2023.html>

Reviewer for:

AAAI 2026, NeurIPS 2025, ICCV 2025, ICML 2025, CVPR 2025, ICLR 2025, AAAI 2025, NeurIPS 2024, ICML 2024, ICLR 2024, AAAI 2024, NeurIPS 2023, AISTATS 2023, AAAI 2023, ICLR 2023, NeurIPS 2022, ICML 2022, AISTATS 2022, AAAI 2022, ICLR 2022, NeurIPS 2021, ICML 2021, AAAI 2021, NeurIPS 2020

Teaching

2020 Spring, TA in **Big Data and Machine Intelligence**, instructed by Zhen Chen

2019 Fall, TA in **Big Data and Machine Intelligence**, instructed by Zhen Chen

2019 Spring, TA in **Machine Learning**, instructed by Prof. Jun Zhu

PS

- OpenRL Lab: <https://github.com/OpenRL-Lab>
- Personal Website: <https://huangshiyu13.github.io>
- 知乎: <https://www.zhihu.com/people/huangshiyu.me>
- GitHub: <https://github.com/huangshiyu13>
- Google Scholar: <https://scholar.google.com/citations?user=PK57vrQAAAAJ&hl=en>
- Linkedin: <https://www.linkedin.com/in/shiyu-huang-841b92106/>