

JIAYI WENG

(412) 996-7548 • trinkle23897@gmail.com • GitHub (3.8k followers) • Google Scholar • LinkedIn

EXPERIENCE

OpenAI L.P.

Member of Technical Staff

San Francisco, CA

07/2022 – now

Public works:

- ChatGPT: initial release effort, 6th author in the contribution list. <https://openai.com/index/chatgpt/>
- GPT-4: RL infra author. <https://openai.com/contributions/gpt-4/>
- GPT-4V: multimodal RL.
- GPT-4 Turbo: post-training.
- GPT-4o: post-training infra lead. <https://openai.com/gpt-4o-contributions/>
- o-series model: post-training infra. <https://openai.com/openai-o1-contributions/>
- Operator: early RL effort. <https://openai.com/index/introducing-operator/>
- GPT-4.5: post-training. <https://openai.com/index/introducing-gpt-4-5/>
- RFT (reinforcement fine-tuning): core RL infra author.
- GPT-5: post-training infra.

EDUCATION

Carnegie Mellon University, School of Computer Science

Pittsburgh, PA

Master in Computational Data Science (System Track), QPA: 3.83/4

08/2020 – 05/2022

Tsinghua University

Beijing, China

B.Eng. in Computer Science and Technology, GPA: 3.77/4, Rank: 18/154

08/2016 – 07/2020

SELECTED OPEN-SOURCE PROJECTS

⌚ Tianshou: A PyTorch Deep Reinforcement Learning Library, ☆ 8686

03/2020 – 08/2020

- Implemented 20+ RL algorithms, imitation/offline RL and multi-agent RL within only 4000 lines of code
- Supported any type of environment state, RNN-style training and parallel environment simulation (synchronous or asynchronous) for all algorithms with complete documentation and comprehensive unit tests
- Received over 8k GitHub stars with more than 50 contributors around the world

⌚ EnvPool: A Highly Parallel RL Environment Execution Engine, ☆ 1183

05/2021 – 12/2021

- Implemented a high-performance parallel environment executor *EnvPool* with C++11 threadpool and pybind11 for a faster reinforcement learning agent training pipeline, about 20x faster than Python implementation and 2x faster than state-of-the-art solution

PERSONAL PUBLICATIONS

- **Jiayi Weng**, Min Lin, Shengyi Huang, Bo Liu, Denys Makoviichuk, Viktor Makoviychuk, Zichen Liu, Yufan Song, Ting Luo, Yukun Jiang, Zhongwen Xu, and Shuicheng Yan, “EnvPool: A Highly Parallel Reinforcement Learning Environment Execution Engine”, in **NeurIPS’22 Datasets and Benchmarks** [PDF]
- **Jiayi Weng**, Huayu Chen, Dong Yan, Kaichao You, Alexis Duburcq, Minghao Zhang, Yi Su, Hang Su, and Jun Zhu, “Tianshou: A Highly Modularized Deep Reinforcement Learning Library”, in **JMLR MLOSS**, [PDF]
- Dong Yan, **Jiayi Weng**, Shiyu huang Chongxuan Li, Yichi Zhou, Hang Su, and Jun Zhu, “Deep Reinforcement Learning with Credit Assignment for Combinatorial Optimization”, in **Pattern Recognition**, [PDF]
- Shihong Song*, **Jiayi Weng***, Hang Su, Dong Yan, Haosheng Zou, and Jun Zhu, “Playing FPS Game with Environment-aware Hierarchical Reinforcement Learning”, in **IJCAI’19 (oral)**, [PDF]
- **Jiayi Weng**, Tsung-Yi Ho, Weiqing Ji, Peng Liu, Mengdi Bao, and Hailong Yao, “URBER: Ultrafast Rule-Based Escape Routing Method for Large-Scale Sample Delivery Biochips”, in **TCAD’18**, [PDF]