Yiran Wu

PA 16803

ykw5399@psu.edu

+1 814-699-1237

RESEARCH INTERESTS

Large Language Models, Multi-agent Systems, Reinforcement Learning

EDUCATION

PhD in Informatics ***Aug 2022 to date***

MSc in Informatics ***Aug 2022 – Dec 2024***

Pennsylvania State University, State College, PA, USA

MSc in ML & AI ***Oct 2021 - Sep 2022***

Imperial College, London, England Graduated with Distinction

BSc in Computer Science ***Sep 2017 - Aug 2020***

University of California, Davis, CA, USA GPA: 3.86/4.0

Dean's Honor List in Spring 2020, Winter 2020, Spring 2019, Spring 2018, Fall 2017

RELATED EXPERIENCES

Research Intern @ Microsoft Research **Redmond, WA | May 27th 2025 to date**

- **Towards General Purposed SWE-agent**: Build agentic scaffold and perform Agentic RL training on LLMs for general-purposed Software Engineering (SWE) Agents.

Research Intern @ Microsoft Research **Redmond, WA | Jun 3rd - Aug 26th, 2024**

- **Benchmarking LLM Agents in Cybersecurity Investigation [publication]**: Build the first benchmark to test LLM-based agents on threat investigation in the form of security question-answering pairs, including a new dataset and a MySQL environment. Propose a new way to generate questions from a bipartite graph with LLMs, enabling a fine-grained step-wise evaluation of agents.

Research Assistant @ Pennsylvania State University **State College, PA | Aug 2022 to date**

- **SimpleDoc** [publication]: (Project Leader) Propose a method to retrieve document pages by first retrieving candidates through embedding similarity and then filtering and re-ranking these candidates based on page summaries, which significantly improves retrieval efficiency and outperform previous methods by 3.2% on 4 document visual question answering tasks.
- **Absolute Zero Reasoner** [publication]: (Second Author, substantial contribution) Propose a new RLVR paradigm called Absolute Zero, in which a single model learns to propose tasks that maximize its own learning progress and improves reasoning by solving them, without relying on any external data. Build and train Absolute Zero Reasoner (AZR), which is a system that self-evolves its training curriculum and reasoning ability by using a code executor to both validate proposed code reasoning tasks and verify answer.
- **AutoGen** [publication]: Co-creator and maintainer of AutoGen, an open-source LLM framework recognized with 45k GitHub stars). Supported key features (function_call, math_proxy_agent, compression, etc) to extend LLM's capabilities. Actively maintaining the repo to ensure robust performance and up-to-date features.
- **Enhance LLM Task-Solving through State-Driven Workflows** [publication]: Proposed StateFlow, a framework that conceptualizes LLM workflows as state machines. StateFlow achieves success rates up to 13% and 28% higher compared to competitive baselines in InterCode SQL and ALFWorld benchmark, with up to 5× and 3× less cost.

- **Defense Again LLM jailbreak attacks** [publication]: Proposed AutoDefense, a multi-agent defense framework against jailbreak attacks. We employ a response-filtering mechanism that has minimal effect on the model's output, while achieve 13% lower Attack Success Rate (ASR) than pervious methods.
- **Solve Math Problems with GPT-4** [publication]: Evaluated the ability of GPT-4 to solve challenging math problems with different methods (Vanilla, PoT, PS) and proposed a novel conversational problem-solving framework *MathChat* that improves over previous methods by 6% in accuracy.
- **Off-policy Learning to Rank with Reinforcement Learning** [publication]: Proposed to use self-attention layer to encode state representations and leveraged several methods to learn state representations.
- **Hyperparameter Tuning with temporal shifted data** [publication]: Proposed validation dataset construction based on chronologically order, which is an important factor for the good performance on temporal data.

Research Assistant @ Sato Lab **Davis, CA (+Remote) | May 2020 - Feb 2022**

- **Classification of heart pixels in cardiac movie data** [publication]: Devised a hybrid unsupervised and supervised machine learning strategy to accurately segment cardiac images, enhancing statistical analysis reliability without manual labeling.

Machine Learning Engineer @ Apulis Tech Inc **Shenzhen, China | Mar 2021 - Aug 2021**

- Supported end-to-end inference and training of one-shot object detection method SSD on Apulis' ML platform.
- Collaborated in the design of a CV Auto-DL pipeline and lead the building of the dataset similarity module.
- Collaborated in PCB defect detection project, leveraging several self-supervised learning (SimCLR, Moco, etc.) to utilize unlabeled data and improve the quality of the backbone. Built a student-teacher network from "Student Teacher Feature Pyramid Neural Networks" for the task.

Machine Learning Intern @ National Supercomputing Center **Wuxi, China | Oct 2020 - Mar 2021**

- Converted and validated a key TensorFlow-based model to PyTorch, aligning with the research from *'Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations' (Science Issue 6481).* Developed an enhanced time-series model using conv-LSTM to advance the application of machine learning in hydromechanics.

PUBLICATIONS

Yiran Wu, Mauricio Velazco, Andrew Zhao, Manuel Raúl Meléndez Luján, Srisuma Movva, Yogesh K Roy, Quang Nguyen, Roberto Rodriguez, Qingyun Wu, Michael Albada, Julia Kiseleva, Anand Mudgerikar. (2025). ExCyTIn-Bench: Evaluating LLM agents on Cyber Threat Investigation. arXiv preprint arXiv:2507.14201.

Chelsi Jain*, **Yiran Wu***, Yifan Zeng*, Jiale Liu, Zhenwen Shao, Qingyun Wu, Huazheng Wang. (2025). SimpleDoc: Multi-Modal Document Understanding with Dual-Cue Page Retrieval and Iterative Refinement. EMNLP 2025.

Andrew Zhao, **Yiran Wu**, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, Gao Huang. "Absolute Zero: Reinforced Self-play reasoning with zero data". Neurips 2025 Spotlight.

Yiran Wu, Tianwei Yue, Shaokun Zhang, Chi Wang and Qingyun Wu. 2024. "StateFlow: Enhancing LLM Task-Solving through State-Driven Workflows". Conference on Language Modeling 2024.

Yifan Zeng*, **Yiran Wu***, Xiao Zhang, Huazheng Wang, Qingyun Wu. 2024. "AutoDefense: Multi-Agent LLM Defense against Jailbreak Attacks". NeurIPS 2024 Workshop on Safe Generative AI

Qingyun Wu, Gagan Bansal, Jieyu Zhang, **Yiran Wu**, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2023. "AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation". Conference on Language Modeling 2024.

Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2023. "MathChat: Converse to Tackle Challenging Math Problems with LLM Agents". ICLR 2024 Workshop on LLMAgents.

Zeyu Zhang, Yi Su, Hui Yuan, **Yiran Wu**, Rishab Balasubramanian, Qingyun Wu, Huazheng Wang, and Mengdi Wang. 2023. "Unified Off-Policy Learning to Rank: a Reinforcement Learning Perspective". In Proceedings of NeurIPS 2023.

Shaokun Zhang, **Yiran Wu**, Zhonghua Zheng, Qingyun Wu, and Chi Wang. 2023. "HyperTime: Hyperparameter Optimization for Combating Temporal Distribution Shifts". ACM MM 2024.

Yiran Wu, Zhen Wang, Crystal M. Ripplinger, and Daisuke Sato. 2022. "Automated Object Detection in Experimental Data Using Combination of Unsupervised and Supervised Methods". In Frontiers in Physiology: Nonlinear Analysis and Machine Learning in Cardiology.

Links

- GitHub Profile
- Personal Website
- Google Scholar

SCHOLARSHIPS

- UCD Annual Fund-SCH 2019
- Pepsi Non-Athletic Scholarship 2019