YIRAN WU

ykw5399@psu.edu Google Scholar Website GitHub

RESEARCH INTERESTS

Large Language Models, Multi-agent Systems, Reinforcement Learning

EDUCATION

PhD in Informatics Aug 2022 – Present
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 (+Remote) | May 27th 2025 – Present

- Build agentic scaffold with LLMs for general-purposed Software Engineering (SWE) Agents. Conducting multi-node **agentic reinforcement learning** and environment scaling with **Kubernetes**.

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

- **ExCyTIn-Bench** [Link]: the first benchmark for **LLM-based agents in cyber threat investigation**, built as security QA pairs with a MySQL environment. Developed a **bipartite-graph-based question generation** method. Benchmark is being integrated into Microsoft Security Solutions.

Research Assistant @ Pennsylvania State University State College, PA | Aug 2022 – Present

- **AgentVM** [Link]: Proposed an **agent-native computer** unifying **GUI + text tools** over a **shared OS state**; improved OSWorld via staged interface augmentations, and demonstrated on GDPVal case studies that it enables **end-to-end tasks** prior environments cannot reliably complete.
- **Live-Evo** [Link]: Proposed an **online self-evolving memory** framework for LLM agents on streaming tasks; over a 10-week live horizon, improves **market returns by 12.9%** (**$408 vs $247**). Live-Evo powered AG2 agent achieved **top 1** on market return on Prophet Arena for 5 weeks.
- **SimpleDoc** [Link]: Proposed a dual-stage document page retrieval method (embedding-based candidate retrieval + summary-based re-ranking), improving performance by **3.2%** on **4 DocVQA tasks**.
- **Absolute Zero Reasoner** [Link]: Proposed **Absolute Zero**, a **self-evolving RLVR** paradigm where a single model **generates and learns from tasks** without external data; built and trained **AZR**, using a code executor to validate tasks and verify solutions for curriculum self-evolution.
- **AutoGen** [Link]: Core creator and maintainer of open-source multi-agent LLM framework AutoGen (**52k+ GitHub stars**), implementing key features such as function calling, conversation compression.
- **StateFlow** [Link]: Proposed a state-machine–based workflow framework for LLMs, achieving up to **13%–28%** higher success rates on InterCode SQL and ALFWorld with up to **5**× and **3**× lower cost.
- **AutoDefense** [Link]: Proposed a multi-agent defense framework against jailbreak attacks. Employed a response-filtering mechanism that achieves a **13%** lower Attack Success Rate (ASR) than previous methods.
- **MathChat** [Link]: Evaluated GPT-4's ability to solve challenging math problems (Vanilla, PoT, PS) and proposed a conversational problem-solving framework that improves accuracy by **6%** over prior methods.

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

- Developed a **hybrid unsupervised + supervised** pipeline to classify heart pixels in cardiac movie data, enabling accurate segmentation and more reliable statistical analysis without manual labels [Link].

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

- Designed a **computer vision Auto-DL pipeline** and implemented the **dataset similarity** module, and support end-to-end training and inference of **one-shot object detection** (SSD) on Apulis' ML platform.
- Improve models on **PCB defect detection** using **self-supervised learning** (SimCLR, MoCo) and feature pyramid networks to leverage unlabeled data.

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

- Converted and validated a key **TensorFlow** model to **PyTorch** to make it compile on the specific GPU architecture. Designed an enhanced **conv-LSTM time-series model** to advance ML applications in hydromechanics.

PUBLICATIONS

Yiran Wu*, Jiale Liu*, Jieyu Zhang, Yaolun Zhang, Shilong Liu, Chi Wang, Mengdi Wang, Huazheng Wang, Qingyun Wu. (2026). Position: Digital Agents Require Unified Agent-Native Computers. ICML 2026 Position.

Yaolun Zhang*, **Yiran Wu***, Yijiong Yu, Qingyun Wu, Huazheng Wang. (2026). "Live-Evo: Online Evolution of Agentic Memory from Continuous Feedback".

Yiran Wu, M. Velazco, A. Zhao, M. R. Meléndez Luján, S. Movva, Y. K. Roy, Q. Nguyen, R. Rodriguez, Q. Wu, M. Albada, J. Kiseleva, and A. Mudgerikar. (2025). "ExCyTIn-Bench: Evaluating LLM agents on Cyber Threat Investigation". ICML 2026.

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.

Q. Wu, G. Bansal, J. Zhang, **Yiran Wu**, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu, A. H. Awadallah, R. W. White, D. Burger, and C. Wang. 2023. "AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation". Conference on Language Modeling 2024.

Yiran Wu, F. Jia, S. Zhang, H. Li, E. Zhu, Y. Wang, Y. T. Lee, R. Peng, Q. Wu, and C. 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". 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". Frontiers in Physiology.

(***** indicates **equal contribution**.)

SKILLS

Programming & Frameworks	Python, AutoGen, LangChain, VERL, VLLM, PyTorch, SQL
Large Language Models	Multi-Agent Systems, Function Calling, Reasoning, Reinforcement Learning, Jailbreak, RAG
Services	Kubernetes, Docker, Git, AWS, Azure