Kerui Wu

✉ keruiwu@umass.edu • keruiwu.github.io • Google Scholar

Education

University of Massachusetts Amherst (UMass Amherst) Sept. 2025 – Present
Ph.D. in Computer Science, GPA: 3.90/4.0
Advisor: Prof. Mingda Qiao

Rensselaer Polytechnic Institute (RPI) Sept. 2021 – May 2025
B.S. in Computer Science, GPA: 3.76/4.0 *Magna Cum Laude*

Research Focus

- **Learning Systems** (Reinforcement Learning, Online Learning)
- **LLM Agents** (Self-Evolving AI, Sequential Decision-Making, Multi-Agent Coordination, Resource Allocation)

Publications

(* Co-first authors; ($\alpha - \beta$) alphabetic order)

- **Adversarially Batched Online Learning**
  - ($\alpha - \beta$) Mingda Qiao, Kerui Wu
  - Under Submission
- **Compute Allocation for Self-Evolving LLMs: From Depth-Breadth to Multi-Armed Bandits**
  - Kerui Wu*, Sixue Xing*, Haoyu He*, Zhuo Yang, Haozheng Luo, Tianfan Fu, Aarthy Nagarajan
  - **EMNLP 2026** (Acceptance Rate: 15%)
- **ClinicalReTrial: Clinical Trial Redesign with Self-Evolving Agents**
  - Sixue Xing, Kerui Wu, Xuanye Xia, Meng Jiang, Jintai Chen, Tianfan Fu
  - **EMNLP 2026** , Findings (Acceptance Rate: 29%)
- **On the Adversarial Robustness of Graph Neural Networks with Graph Reduction**
  - Kerui Wu, Ka-Ho Chow, Wenqi Wei, Lei Yu
  - **ESORICS 2025** (Acceptance Rate: 16%) — **Oral**
- **Large EEG-U-Transformer for Time-Step level Detection**
  - Kerui Wu, Ziyue Zhao, Bülent Yener
  - Technical Report, **arXiv** — **Winning Model of 2025 Seizure Detection Challenge**
- **On the Robustness of Graph Reduction Against GNN Backdoor Attacks**
  - Yuxuan Zhu, Michael Mandulak, Kerui Wu, George Slota, Yuseok Jeon, Ka-Ho Chow, Lei Yu
  - **AISec@CCS 2024** (Acceptance Rate: 25%)

Awards

NSF NAIRR Pilot Start-Up Award, PI 2026
National Artificial Intelligence Research Resource (NAIRR)
CICS Scholarship 2025
Manning College of Information & Computer Sciences, UMass Amherst
Best Algorithm Award 2025
2025 Seizure Detection Challenge, International Conference on AI in Epilepsy
Certified in CyberSecurity (CC) 2023
ISC^2 (International Information System Security Certification Consortium)

Research Experience

Self-Evolving Multi-Agent Systems Sept. 2025 – Present

- Modeling **computation resource allocation** for the **LLM-based self-evolving process** to balance **exploration-exploitation** trade-offs.
- Utilize **Multi-Armed Bandits** algorithms to adaptively allocate resources across evolving trajectories, achieving 12.3% improvement on three **QWen3** and three **Llama** models over three tasks.
- Develop *ClinicalReTrial*, a reasoning-driven **LLM-based multi-agent framework** that formulates clinical trial redesign as an iterative protocol optimization problem.
- Integrate **retrieval-augmented generation (RAG)** and **hierarchical memory with in-context learning** to enable continual agent self-improvement, achieving 83.3% protocol improvement and 5.7% success probability gain.

Provable Online & Collaborative Learning Sept. 2025 – Present

- Develop theoretical foundations for **adversarially batched online learning**, where the model may update its policy only at a restricted set of time steps.
- Design learning algorithms under adversarially batched settings that achieve **provable optimal performance**.
- Design an **generic adaptive doubling** algorithm for learning algorithms to run without knowing the updating time points while preserving optimal performance guarantees.
- Design collaborative learning algorithms under natural and verifiable conditions with mathematically provable efficiency and **low sample complexity**.

Representation Learning for Bio-Signal May 2024 – May 2025

- Propose a time-step-level **sequence modeling paradigm** for long-sequence EEG analysis, eliminating redundant sliding-window inference and enabling **event-centric classification**.
- Design a scalable **Conv-Transformer hybrid architecture** integrating multi-scale CNN encoders, residual convolution stacks, and **self-attention based temporal embedding modules**.
- Achieve 10× runtime speedup while reaching 19.4% F1-score improvement, winning **1st Place** in the 2025 Seizure Detection Challenge (International Conference on AI in Epilepsy).

Trustworthy ML Against Adversarial Attacks Jan. 2024 – May 2025

- Conduct a systematic robustness study of **Graph Neural Networks** under adversarial settings, including **poisoning, evasion, and backdoor attacks**.
- Investigate the security implications of **graph reduction algorithms** in scalable GNN systems, revealing how structural graph transformations interact with adversarial perturbations.

Selected Projects

China Construction Bank Achievement System React, Go, MySQL, Azure

- Developed a production web platform for tracking employee performance, **deployed across 11 branches of China Construction Bank.**
- Implemented full-stack functionality, including data dashboards, analytics queries, and role-based management tools.
- Designed secure authentication and API infrastructure using **SHA256 hashing and JWT**, and deployed the system on **Microsoft Azure cloud**.

Submitty Open-Source Contribution PHP, JavaScript, SQL, Python

- Contributed to Submitty, an open-source **course management** and **automated grading** system adopted by multiple universities.
- Implemented **full-stack features** including customizable pronoun settings across frontend UI, backend APIs (PHP), and relational database schemas (SQL). **5 pull requests were merged into the main branch of the repository.**

Technical Skills

Machine Learning: PyTorch, Torch-Geometric, Scikit-Learn, vLLM, LangChain, MCP
Programming: Python, C/C++, JavaScript, Go, PHP, SQL