Yu Wu

☎551-229-5305 ✉yw828@rutgers.edu https://yuwuofrutgers.github.io

RESEARCH INTERESTS

• **Efficient and reliable AI for hybrid system under resource constraints**: Learning with reject option, Learning to Defer, Selective Prediction, Reinforcement Learning, distributed optimization, early-exiting
• **LLMs Collaboration**: LLM Routing, LLM Cascade, In-context Learning, multi-agent systems, model orchestration, memory systems, human–computer interaction, VLM

EDUCATION

Rutgers, The State University of New Jersey, New Jersey, USA *Aug. 2020 - present*
PhD candidate in Electrical and Computer Engineering (Advisor: Prof. Anand D. Sarwate)
University of Science and Technology of China (USTC), Anhui, China *Jun. 2017 - Jun. 2020*
Master's in Electronic Engineering and Information Science (Advisor: Prof. Bin Liu)
University of Science and Technology of China (USTC), Anhui, China *Jul. 2013 - Jun. 2017*
Bachelor's in Information Security

SELECTED RESEARCH PROJECTS

• **Interactive LLM Cascade** | LLM Routing, RAG, In-context Learning, model orchestration, agent system, Distillation, memory system
*- **Project Leader**, Research Assistant at Rutgers* *Sep. 2024 - Sep. 2025*
*- To achieve higher accuracy and save tokens in **multi-LLM collaboration**, we propose Interactive-LLM-Cascade, an online **RAG**-augmented system where weaker LLM can not only defer to stronger LLM for difficult questions, but also gets boosted by stronger LLM's feedback in a long term.*
*- Inter-Cascade improves the accuracy of the weaker model (by up to **33.06%)** and the overall system (by up to **5.53%**), while reducing the calls to strong models (by up to **48.05%**) and saving the corresponding fees (by up to **49.63%**).* [paper]

• **Collaborative Machine Learning for Client-server System**| model routing, real-time inference, distributed system, CNN, ViT
*- **Project Leader**, Research Assistant at Rutgers* *Sep. 2020 - Sep. 2024*
- To enhance legacy ML models which are resource-constrained, we propose the Learning to Help (L2H) diagram to jointly train external machine/human experts, and a query router that efficiently guides data samples to suitable models.
*- L2H is theoretically **Bayes optimal** and increases the system overall accuracy by **4%~12%** and works for different **distributed** scenarios.* [paper1, paper2, code]

• **Reasoning Teacher–Student Framework for End-to-End Autonomous Driving**| Reasoning VLM, Distillation, GRU
- Project collaborator *Apr. 2025 - Nov. 2025*
- To improve the interpretability and robustness of end-to-end autonomous driving under the absence of explicit reason–action annotations, we propose a teacher–student architecture in which the teacher model generates refined reasoning, and the student model distills this knowledge and predict numerical trajectory.
*- Our method achieves a **24%** performance improvement over non-reasoning baselines, while providing enhanced **interpretability** and maintaining suitability for deployment on **resource-constrained devices**.*

• **Enhancing Model-Based Reinforcement Learning with Data Filter**| Out-of-distribution, RL, MuJoCo, MBPO, Actor-Critic
- Project collaborator *Jun. 2024 - Sep. 2024*
*- To efficiently bridge model-free and model-based Reinforcement Learning, we propose **Out-of-distribution** (OOD) data filter to **Dyna-Style Model-Based RL**, which adaptively filters the deviated data generated by estimated model.*

*-Our theorems provide **tighter bound** for estimated error and our experiments save up to **25%** epochs to reach convergence. [paper]*

• **Human–computer Interactive Sensing** I wearable device, human-centered computing, human-computer interaction

- Researcher (Intern) at X-discovery Lab, Dartmouth College *Dec. 2018 - Apr. 2019*

*- To achieve human-computer interaction through soft material, we propose an inductive sensing based prototype called Tessutivo. We yielded **93.9% real-time accuracy** for object recognition.*[paper, demo]

• **Anti-interference for WiFi-based Human Activity Recognition (HAR)**ICSI, non-intrusive sensing, PyTorch

- Research Assistant at EEIS department of USTC *Sep. 2017 - May 2020*

*- To mitigate the interference components in WiFi signals, we propose PrhaseAnti system. Our method **improves up to 16%** on accuracy and **9**× **faster** recognition speed.*[paper1, paper2]

SELECTED WORK EXPERIENCE

• **Online Learning Algorithm for Audio Echo Cancellation** I least mean square, adaptive filtering, acoustic signal processing

*- Researcher (Intern) at **Nokia Bell Lab*** *Jun. 2022 - Aug. 2022*

- Toward the misalignment error of online estimation system, we analyze the accuracy and effectiveness of "delay and extrapolate" algorithm, which is widely used in adaptive filter.

*- Our method proves the **optimality** and **corrects misuse** in previous works. The experimental results are **consistent** with the theoretical derivation.* [paper]

PUBLICATIONS

• **Y Wu**, S Wu, et al., "Not only a helper, but also a teacher: Interactive LLM Cascade." *Under review by ICLR 2026*

• **Y Wu**, Y Li, et al., "Learning to Help in Multi-Class Settings." *ICLR 2025*

• **Y Wu**, and Anand Sarwate, "Learning to Help: Training Models to Assist Legacy Devices." *ISIT 2024 Workshop IT-TML*

• Z Dong, Y Zhu, **Y Wu**, et al., "FROST-Drive: Scalable and Efficient End-to-End Driving with a Frozen Vision Encoder." *WACV 2026 Workshop LLVM-AD*

• J Huang, B Liu, C Miao, Y Lu, **Y Wu**, et al., "PhaseAnti: An anti-interference WiFi-based activity recognition system using interference-independent phase component." *IEEE Transactions on Mobile Computing 2021*

• J Huang, B Liu, P Liu, C Chen, N Xiao, **Y Wu**, et al., "Towards anti-interference WiFi-based activity recognition system using interference-independent phase component." *INFOCOM 2020*

• J Gong, **Yu Wu,** et al., "Tessutivo: Contextual interactions on interactive fabrics with inductive sensing." *UIST 2019*

• W Zhang, Y Li, Z Dong, **Y Wu**, et al., "Renaissance of Literate Programming in the Era of LLMs: Enhancing LLM-Based Code Generation in Large-Scale Projects." *arXiv (2024)*

• Z Yu, L An, Y Li, **Y Wu**, et al., " EAPCR: A Universal Feature Extractor for Scientific Data without Explicit Feature Relation Patterns." *arXiv (2024)*

• Y Li, Z Dong, E Luo, **Y Wu**, et al., " When to trust your data: enhancing dyna-style model-based reinforcement learning with data filter." *arXiv (2024)*

SKILLS

Programming: Python, C, C++, MATLAB, Processing, SQL, Java, R

Technical tools: PyTorch, NumPy, SciPy, Pandas, scikit-learn, Slurm, Git, LoRA, HuggingFace

Course: Qishi Deep Learning in Recommender Systems Bootcamp

SERVICE

Reviewer: ICLR, IJCNLP&AACL, TMLR, Journal of Knowledge-Based Systems

Teaching Assistant: Probability and Random Processes, Linear Signal and System, Discrete Mathematics, etc.