

Ruhan Wang

📍 Bloomington, IN, USA ✉️ ruhwang@iu.edu 📞 +1 (812) 327-8256 🔗 [LinkedIn](#) 🎓 [Google Scholar](#)

- **Indiana University** Aug. 2022 – Present
Ph.D. of Computer Engineering *Bloomington, IN, US*
- **Indiana University** Aug. 2022 – Dec. 2024
Master of Computer Engineering *Bloomington, IN, US*

RESEARCH INTERESTS

- Agent Reinforcement Learning
- Reinforcement Learning for Post-training, Reasoning
- Self-Improving Large Language Models

HONORS AND AWARDS

- First Prize, Contemporary Undergraduate Mathematical Contest in Modeling, 2021
- National Excellent Award, The Contest of Lan Qiao Cup, 11/2020
- First Prize & Second Prize, Contemporary Undergraduate Mathematical Contest in Modeling, 11/2020
- Scholarship of Henan University, 2020 & 2019
- First Prize, Chinese Mathematics Competition, 2019
- Outstanding Graduates of Henan University, 2022
- Excellent thesis in Henan Province, China, 2022

WORKING EXPERIENCE

- Ph.D. Research Intern at Tencent AI Lab (Hunyuan Frontier Lab), Bellevue, WA, May 2026 - Aug. 2026 (Host: [Kishan Panaganti](#)). Working on Agentic Reinforcement Learning for large language model agents.
- Ph.D. Research Intern at Mitsubishi Electric Research Laboratories, Cambridge, MA, May 2024 - Aug. 2024 (Host: [Toshiaki Koike-Akino](#)). Worked on quantum machine learning and generative models.
- Graduate Student Researcher, Machine Learning Lab, Indiana University, Aug. 2023 - Present (Advisor: [Dr. Dongruo Zhou](#)).
- Graduate Student Researcher, Quantum Computing Lab, Indiana University, Aug. 2022 - Aug. 2023 (Advisor: Dr. Fan Chen).

TEACHING EXPERIENCE

- Associate Instructor for CSCI-B659: Reinforcement Learning for LLM. (Spring 2026)
- Associate Instructor for CSCI-B565: Data Mining. (Fall 2024 & Fall 2025)
- Associate Instructor for CSCI-B455: Principles of Machine Learning. (Spring 2025)
- Associate Instructor for ENGR-E516: Engineering Cloud Computing. (Spring 2024)
- Associate Instructor for ENGR-E599: Deep Learning Architecture and Hardware Acceleration. (Fall 2023)

PROFESSIONAL SERVICES

- Conference Reviewer: International Conference on Learning Representations (ICLR) — **Top 25% Reviewer, 2026**
- Conference Reviewer: International Conference on Machine Learning (ICML) — **Gold Reviewer, 2026**
- Conference Reviewer: Annual Conference on Neural Information Processing Systems (NeurIPS)
- Conference Reviewer: Association for the Advancement of Artificial Intelligence (AAAI)
- Journal Reviewer: Measurement Science and Technology
- Journal Reviewer: Physica Scripta

RESEARCH EXPERIENCE

- **More Memory, Worse Agents: Error Reproduction and Anti-Persistence in LLM Agents** Indiana, US
Hosts: Kishan Panaganti (Tencent AI Lab), Assistant Prof. Dongruo Zhou (Indiana University) Feb. 2026 – Present
 - Identified error reproduction as a structural failure mode in persistent abstraction-based memory for self-improving LLM agents, distinct from trajectory-level error propagation.
 - Proposed ANTI-PERSISTENCE, a memory module that stores only factual interaction records (raw trajectories and outcomes) while constructing task-specific abstractions on demand at inference time, removing the cross-iteration feedback loop induced by persistent abstractions.
 - Designed an adaptive mode-selection mechanism over four abstraction modes (none, imitation, strategic, corrective) via a recency-weighted contextual UCB bandit operating on a factual outcome log, enabling reusable abstraction procedures without persisting abstraction contents.
 - Conducted comprehensive experiments across three benchmarks (ALFWorld, WebShop, ScienceWorld) and three backbones (Llama-4-Scout, Qwen3-Coder-Next, Gemma-4-31B-it), demonstrating consistent improvements in held-out task success rate, inference cost, and long-term iterative stability over persistent-memory baselines.
- **Uncertainty-Aware Federated Reasoning with Large Language Models (FERA)** Indiana, US
Advisor: Assistant Prof. Dongruo Zhou Aug. 2025 – Feb. 2026
 - Proposed FERA, a parameter-free federated reasoning framework that leverages uncertainty quantification to address the trade-off between computational efficiency and reasoning performance in federated rLLMs.
 - Developed a dual-pipeline aggregation mechanism: uncertainty-weighted majority voting for simple QA tasks and uncertainty-aware self-critique for complex Chain-of-Thought reasoning.
 - Designed iterative server–client collaboration protocols that progressively refine uncertainty estimates across communication rounds.
 - Conducted comprehensive experiments on diverse benchmarks, demonstrating consistent improvements over training-based and parameter-free baselines.
- **Agentic Recommender Systems with Multimodal LLMs** Indiana, US
Advisor: Assistant Prof. Dongruo Zhou Dec. 2024 – Mar. 2025
 - Develop a formal framework for LLM-based Agentic Recommender Systems (LLM-ARS), incorporating core modules such as user profiling, memory, planning, and action selection to enable continuous learning and proactive recommendation generation.
 - Explore agentic capabilities of Large Language Models through multi-agent collaboration and role-playing simulations; benchmark planning and reasoning abilities using chain-of-thought prompting, user feedback modeling, and surrogate user environments.
 - Identify and analyze seven key research challenges in LLM-ARS, including multimodal fusion, interaction grounding, lifelong personalization, and safe autonomous behavior; propose future directions for scalable, interpretable, and ethically-aligned RS systems.
- **Federated In-context Learning: Make Agent Powerful** Indiana, US
Advisor: Assistant Prof. Dongruo Zhou Oct. 2024 - Jan. 2025
 - Propose a novel privacy-preserving framework for Federated In-Context LLM Agents(Fed-ICL), leveraging the strengths of in-context learning to collaboratively train diverse LLM agents while ensuring data privacy through federated learning.
 - Establish a theoretical foundation for Fed-ICL by proving its equivalence to established federated learning algorithms, highlighting the robustness and effectiveness of the proposed framework.
 - Compare the proposed framework with other collaborative LLM agent baselines, such as Mixture-of-Agents and Multiagents Debate, across various tasks. Experimental results demonstrate the superior performance of our algorithm.
- **Incorporating RL into Retrieval-Augmented Generation (RAG) for Personalizing LLM Agents** Indiana, US
Advisor: Assistant Prof. Dongruo Zhou May 2024 - Aug. 2024
 - Incorporate sequential decision-making into RAG overcomes its limitations, such as the inability to utilize user feedback and generate personalized retrievals, thereby enabling LLM agents to produce tailored content.
 - Propose a contextual bandit-based retrieval framework to optimize retrieval models, enabling the delivery of a personalized selection of documents to LLM agents for enhanced, tailored content generation.
 - Compare the performance of the contextual bandit-based retriever with traditional retrieval methods, such as BM25, highlighting the superior effectiveness of the contextual bandit-based approach in delivering personalized and relevant documents for large language models.
- **Enhance Offline RL with Return Augmentation for Domain Adaptation Problems** Indiana, US
Advisor: Assistant Prof. Dongruo Zhou Feb. 2024 - May 2024
 - Study offline off-dynamics RL to leverage data from an easily accessible source domain, enhancing policy learning and enabling the trained agent to perform effectively in the target domain.
 - Introduce the novel Return Augmented Decision Transformer (RADT) method to address the off-dynamics challenge by augmenting the returns of offline trajectories in the source domain. A rigorous analysis is provided to demonstrate that the return-conditioned policy learned via RADT achieves a comparable level of suboptimality to a policy trained directly in the target domain.

- Conduct comprehensive experiments on the D4RL benchmark under various off-dynamics shift settings to evaluate the performance of RADT, comparing it against other baseline methods and demonstrating its effectiveness.

• **Enhance Safe Offline RL with Constrained Q-learning Decision Transformer**

Indiana, US

Advisor: Assistant Prof. Dongruo Zhou

May 2024 - Aug. 2024

- Introduce the Constrained Q-learning Decision Transformer (CQDT) to address the stitching limitations of the Constrained Decision Transformer (CDT) in the safe reinforcement learning domain. CQDT enables agents to select the optimal policy that maximizes rewards while strictly adhering to safety constraints.
- Compare the performance of CQDT with other baselines across various safe reinforcement learning benchmarks. Experimental results demonstrate the superior effectiveness of CQDT in achieving high rewards while adhering to safety constraints.

• **Model the End-to-end Carbon Footprint of Large Language Models**

Indiana, US

Advisor: Assistant Prof. Fan Chen

Jul. 2023 - Sep. 2023

- Propose an end-to-end carbon footprint projection model, LLMCarbon, designed to accurately predict the carbon footprint of various LLM agents across their training, inference, experimentation, and storage phases.
- Integrate key parameters from LLMs, hardware, and data centers—such as parameter count, hardware specifications, system power consumption, chip area, and data center efficiency—to effectively model the operational and embodied carbon footprints of LLM agents.
- Provide a detailed comparison of various LLM frameworks, leveraging LLMCarbon to compute the carbon footprints associated with different LLM agents.

PUBLICATION

-
- **Ruhan Wang**, Kishan Panaganti, Dongruo Zhou, "More Memory, Worse Agents: Error Reproduction and Anti-Persistence in LLM Agents", *The Fortieth Conference on Neural Information Processing Systems (NeurIPS, Under Review)*, 2026
 - **Ruhan Wang**, Chengkai Huang, Zhiyong Wang, Rui Wang, Tong Yu, Lina Yao, Dongruo Zhou, "Uncertainty-Aware Federated Reasoning with Large Language Models", *Conference on Language Modeling (COLM, Under Review)*, 2026
 - Runze Zhao, Yue Yu, **Ruhan Wang**, Chunfeng Huang, Dongruo Zhou, "[Instance-Dependent Continuous-Time Reinforcement Learning via Maximum Likelihood Estimation](#)", *Forty-Third International Conference on Machine Learning (ICML)*, 2026
 - **Ruhan Wang**, Zhiyong Wang, Chengkai Huang, Rui Wang, Tong Yu, Lina Yao, John C.S. Lui, Dongruo Zhou, "[Federated In-Context Learning: Iterative Refinement for Improved Answer Quality](#)", *Forty-Second International Conference on Machine Learning (ICML)*, 2025
 - Chengkai Huang, Junda Wu, Yu Xia, Sheldon Yu, **Ruhan Wang**, Tong Yu, Ruiyi Zhang, Ryan Rossi, Branislav Kveton, Dongruo Zhou, Julian McAuley, Lina Yao, "[Towards Agentic Recommender Systems in the Era of Multimodal Large Language Models](#)", *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2025
 - Zhishuai Liu, Yu Yang, **Ruhan Wang**, Pan Xu, Dongruo Zhou, "[How to Provably Improve Return Conditioned Supervised Learning?](#)", *arXiv preprint*, 2025
 - **Ruhan Wang**, Ye Wang, Jing Liu, Toshiaki Koike-Akino, "[Quantum Diffusion Models for Few-Shot Learning](#)", 2025 *IEEE International Conference on AI and Data Analytics (ICAD)*, 2025; *AAAI 2024 Quantum Computing and Artificial Intelligence Workshop*, 2024
 - **Ruhan Wang**, Dongruo Zhou, "[Safe Decision Transformer with Learning-based Constraints](#)", *7th Annual Learning for Dynamics and Control Conference (L4DC)*, 2025; *NeurIPS 2024 Safe Generative AI Workshop*, 2024
 - **Ruhan Wang**, Yu Yang, Zhishuai Liu, Dongruo Zhou, Pan Xu, "[Return Augmented Decision Transformer for Off-Dynamics Reinforcement Learning](#)", *Transactions on Machine Learning Research (TMLR)*, 2024
 - Ahmad Faiz, Sotaro Kaneda, **Ruhan Wang**, Rita Osi, Parteek Sharma, Fan Chen, Lei Jiang, "[LLMCarbon: Modeling the End-to-End Carbon Footprint of Large Language Models](#)", *The Twelfth International Conference on Learning Representations (ICLR Oral)*, 2024
 - **Ruhan Wang**, Lei Jiang, Fahiz Baba-Yara, Fan Chen, "[JustQ: Automated Deployment of Fair and Accurate Quantum Neural Networks](#)", *29th Asia and South Pacific Design Automation Conference (ASP-DAC)*, 2024
 - **Ruhan Wang**, Phil Richerme, Fan Chen, "[A Hybrid Quantum-Classical Neural Network for Learning Transferable Visual Representation](#)", *Quantum Science and Technology*, 2023

- **Ruhan Wang**, Ruixin Qiao, Yukang Zou, "[A Brief Analysis on Damaged Building Classification: Optimizer and Learning Rate](#)", *2022 International Conference on Cloud Computing, Performance Computing and Deep Learning (SPIE 12287)*, 2022
- Kaifang Li, Guancheng Hui, **Ruhan Wang**, Miaohui Zhang, "[Person Re-Identification Based on Generative Adversarial Network and Self-Calibrated Convolution](#)", *Laser & Optoelectronics Progress*, 2022

ACADEMIC EXPERIENCE

- **Luddy AI Center's Open House event**: Prepare a presentation for research project "A Hybrid Quantum-Classical Neural Network for Learning Transferable Visual Representation". (Mar. 2023, Bloomington, US)
- **Facial Recognition**: Self-studied OpenCV related image processing, and image noise reduction and image transformation in digital image processing. Used deep learning related algorithms to match face similarity and self-studied deep learning and Pytorch framework. (Mar. 2021, Kaifeng, China)
- **The Realization of Forward Transmission and Reverse Transmission of BP Neural Network**: Studied the process of forward transmission and reverse transmission of BP Neural Network. Mastered different loss functions in deep learning and the characteristics and differences of the optimizer. Used Python to complete the code. (Apr. 2021, Kaifeng, China)
- **Campus Sign-in System Based on Mobile Phone Positioning**: Created tables and conducted operations such as adding, deleting, modifying and checking tables. Deployed the cloud server to manage the deployment of the App backend. (Sep.2020, Kaifeng, China)