

Renhong Huang

Zhejiang University, Hangzhou, China
+86-13620950268 | renh2@zju.edu.cn | renh2.github.io

EDUCATION

Zhejiang University <i>PhD, Computer Science</i> <ul style="list-style-type: none">Advised by Yang Yang	Hangzhou, China 2022.9 - Current
Zhejiang University <i>B.E., Computer Science</i> <ul style="list-style-type: none">GPA: 3.97/4.00	Hangzhou, China 2018.9 - 2022.6
National University of Singapore <i>Visiting Student, Computer Science</i> <ul style="list-style-type: none">Courses: Machine Learning, IoT	Lower Kent Ridge Road, Singapore 2021.8 - 2021.9
Fudan University <i>Visiting Student, Computer Science</i> <ul style="list-style-type: none">Advised by Jiarong Xu	Shanghai, China 2022.9 - 2023.9

PUBLICATIONS

- Renhong Huang**, Ning Tang, Jiarong Xu, Yuxuan Cao, Qingqian Tu, Sheng Guo, Bo Zheng, Yang Yang. *PolicySim: An LLM-Based Agent Social Simulation Sandbox for Proactive Policy Optimization*. WWW 2026 Oral.
- Siwei Zhang, Yun Xiong, Xi Chen, Zi'an Jia, **Renhong Huang**, Jiarong Xu, Jiawei Zhang. *RAPO: Expanding Exploration for LLM Agents via Retrieval-Augmented Policy Optimization*. KDD 2026 Oral.
- Renhong Huang**, Dongdong Hua, Yifei Sun, Sitao Ding, Hanyang Yuan, Daixin Wang, Yang Yang. *RAT: RunAnything via Fully Automated Environment Configuration*. arXiv 2026.
- Dongdong Hua, Yifei Sun, **Renhong Huang**, Feng Gao, Chunping Wang, Yang Yang. *PTCG-Bench: Can LLM Agents Master Pokemon Trading Card Game?*. arXiv 2026.
- Renhong Huang**, Yuxuan Cao, Yi Li, Junwei Hu, Zihua Xiong, Shuai Fang, Sheng Guo, Bo Zheng, Yang Yang. *Trimming the Fat: Redundancy-Aware Acceleration Framework for DGNNs*. AAAI 2026.
- Hanyang Yuan, Ning Tang, Tongya Zheng, Jiarong Xu, Xintong Hu, **Renhong Huang**, Shunyu Liu, Jiacong Hu, Jiawei Chen, Mingli Song. *Tree of Preferences for Diversified Recommendation*. NeurIPS 2025.
- Renhong Huang**, Jiarong Xu, Zhiming Yang, Xiang Si, Xin Jiang, Hanyang Yuan, Chunping Wang, Yang Yang. *Extracting Training Data from Molecular Pre-trained Models*. NeurIPS 2024.
- Hanyang Yuan, Jiarong Xu, **Renhong Huang**, Mingli Song, Chunping Wang, Yang Yang. *Towards More Efficient Property Inference Attacks on Graph Neural Networks*. NeurIPS 2024.
- Renhong Huang**, Jiarong Xu, Xin Jiang, Ruichuan An, Yang Yang. *Can Modifying Data Address Graph Domain Adaptation?* KDD 2024 Oral.
- Renhong Huang**, Jiarong Xu, Xin Jiang, Chenglu Pan, Zhiming Yang, Chunping Wang, Yang Yang. *Measuring Task Similarity and Its Implication in Fine-Tuning Graph Neural Networks*. AAAI 2024.
- Jiarong Xu, **Renhong Huang**, Xin Jiang, Yuxuan Cao, Carl Yang, Chunping Wang, Yang Yang. *Better with Less: A Data-Centric Perspective on Pre-Training Graph Neural Networks*. NeurIPS 2023.
- Yifei Sun, Haoran Deng, Yang Yang, Chunping Wang, Jiarong Xu, **Renhong Huang**, Linfeng Cao, Yang Wang, Lei Chen. *Beyond Homophily: Structure-aware Path Aggregation Graph Neural Network*. IJCAI 2022.
- Jiarong Xu, Zhiming Yang, Xiang Ao, **Renhong Huang**, Chenglu Pan. *Release Graph Pre-Trained Model via Privacy-Preserving Data Augmentation*.

RESEARCH EXPERIENCE

- RAPO: Retrieval-Augmented Policy Optimization for LLM Agents.** 2025.12 - 2026.03
Supervisor: Prof. Jiarong Xu Zhejiang University
- Studied the exploration bottleneck in agentic reinforcement learning, where pure on-policy rollouts often repeat similar reasoning paths and miss useful step-level exploratory behaviors.
 - Developed RAPO, a retrieval-augmented policy optimization framework that injects off-policy step traces into agent rollouts, achieving a 5.0% average performance gain across 14 datasets and 1.2× faster training efficiency.

- RAT: Fully Automated Environment Configuration for LLM Agents.** 2025.12 - 2026.04
Supervisor: Prof. Yang Yang Zhejiang University
- Identified executable environment setup as bottleneck for repository-level code agents, since many agents can write plausible code but fail to install dependencies, configure runtimes, and obtain reliable execution feedback.
 - Proposed RunAnything (RAT), a language-agnostic automated configuration pipeline that improves Environment Setup Success Rate (ESSR) by 29.6% on average over strong baselines, moving agents closer to real deployment.
- PTCG-Bench: Long-Horizon Strategic Benchmarking for LLM Agents.** 2026.02 - 2026.05
Supervisor: Prof. Yang Yang Zhejiang University
- Built PTCG-Bench around Pokemon Trading Card Game to stress-test LLM agents in a single vivid environment requiring hidden-information reasoning, long-horizon planning, stochastic card draws, rule following, and text-numeric decision making.
 - Showed that LLM agents can reach non-trivial gameplay performance, but stable self-evolution from accumulated experience remains difficult and measured performance is highly sensitive to harness design, highlighting a concrete gap between language competence and reliable strategic agency.
- Trimming the Fat: Redundancy-Aware Acceleration Framework for DGNNs.** 2024.12 - 2025.3
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang Zhejiang University & Ant Group
- Temporal graphs are essential for modeling complex real-world systems, but the high computational cost poses practical challenges for deploying dynamic graph neural networks (DGNNs).
 - Proposed TrimDG, a general framework that accelerates DGNNs by eliminating both static and runtime redundancy, reducing runtime by an average of 83.80% across diverse DGNNs.
- LLM-Agent Social Simulation Sandbox for Empowering Intervention.** 2024.12 - 2025.3
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang Zhejiang University & Fudan University
- Studied how social media platforms influence information propagation, where interventions such as recommendation systems may amplify risks if deployed without prior assessment.
 - Proposed PolicySim, an LLM-based social simulation sandbox for proactively evaluating and optimizing platform intervention policies before real-world deployment.
- Tree of Preferences for Diversified Recommendation.** 2025.06 - 2026.01
Supervisor: Prof. Jiarong Xu Zhejiang University & Fudan University
- Diversified recommendation often suffers from data bias, where observed user feedback fails to reflect underexplored preferences, leading to recommendation homogeneity.
 - Proposed Tree of Preferences (ToP), a framework that leverages LLMs' zero-shot inference and world knowledge to uncover and organize underexplored user preferences into a hierarchical tree structure.
- Extracting Training Data from Molecular Pre-trained Models.** 2023.10 - 2024.6
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang Zhejiang University & Fudan University
- Pre-trained models contain a large amount of information. This work provides the first investigation into whether molecular graph pre-trained models carry data extraction risks.
 - Experiments demonstrate that even with only query access to molecular pre-trained models, there is a considerable risk of extracting training data.
- Towards More Efficient Property Inference Attacks on GNN.** 2023.10 - 2024.4
Supervisor: Prof. Jiarong Xu Zhejiang University & Fudan University
- Pre-trained models contain a large amount of information. Privacy leakage issues need to be considered when attackers attempt to extract knowledge (training datasets) from pre-trained models.
 - Proposed an efficient graph property inference attack using model approximation techniques, achieving an average 2.7% improvement in attack accuracy across real-world settings.
- Can Modifying Data Address Graph Domain Adaptation?** 2023.4 - 2023.9
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang Zhejiang University & Fudan University
- Identify limitations of model-only graph domain adaptation and propose two data-centric principles (alignment & rescaling) based on generalization theory.
 - Proposed GraphAlign, a data-centric method that generates a small but transferable graph, outperforming baselines by 3.38% using just 1% of the original graph.
- Measuring Task Similarity in Fine-Tuning Graph Neural Networks** 2022.9 - 2023.1
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang Zhejiang University & Fudan University
- Studies the specific application scope of graph pre-trained models, i.e., the extent to which downstream tasks can benefit from specific pre-training tasks.
 - Introduced task consistency to quantify the alignment between pre-training and downstream tasks, and proposed Bridge-Tune, a fine-tuning strategy that mitigates their gaps.

- Release Graph Pre-trained Model via Privacy-Preserving Data Augmentation.** 2022.8 - 2022.10
Supervisor: Prof. Jiarong Xu *Fudan University*
- Target a privacy-preserving graph pre-training model to facilitate the knowledge sharing environment, which allows data owners to disclose their pre-trained models without exposing privacy.
 - Introduced two causality-inspired principles for graph pre-training: privacy preservation and generalization.
- Better with Less: A Data-Centric Perspective on Pre-Training Graph Neural Networks** 2021.1 - 2022.9
Supervisor: Prof. Jiarong Xu and Prof. Yang Yang *Zhejiang University*
- Conduct extensive experiments to study how pre-training data affects graph pre-training performance, revealing curse of big data where more data does not always yield better results.
 - Proposed a better-with-less framework that uses fewer but carefully selected samples to boost GNN pre-training. Our method surpasses existing graph pre-training baselines by 9.94% (freezing) and 17.83% (fine-tuning).
- Beyond Homophily: Structure-aware Path Aggregation Graph Neural Network** 2021.10 - 2022.1
Supervisor: Prof. Yang Yang *Zhejiang University*
- Shed light on the path level patterns in graphs to generalize GNNs for both homophily and heterophily graphs.
 - Introduced a structure-aware recurrent cell consisting of order-preserving and distance-aware components to learn the semantic information of neighborhoods.

SELECTED AWARDS AND HONORS

National Scholarship, 2024.
Outstanding Graduate Thesis, 2022.
Award of Honor for Graduate at Chu Kochen Honors College of Zhejiang University, 2022.
Third Prize, Zhejiang University's Academic Year Scholarship, 2018 - 2021.
Meritorious Winner for American Mathematical Competition (top 7%), 2021.
Outstanding Academic Model, Zhejiang University, 2018-2021.
Second Prize, Province Mathematics Competition, 2020.

RESEARCH INTERESTS

LLM Agents, LLM Reasoning and Alignment, Humanoid VLA, Environment-Grounded Agents, Data-centric AI.