

Yinjie Wang

yinjie@uchicago.edu

EDUCATION

University of Chicago

2024 - Present

PhD in Statistics

University of Science and Technology of China

2020 - 2024

Bachelor of Science in Probability and Statistics

- GPA: 3.95, **ranked 1st** in school of gifted young majored in Statistics

RESEARCH INTERESTS

Reinforcement learning, diffusion language models, LLMs, agentic AI, and efficient attentions

RESEARCH EXPERIENCE (ALL FIRST AUTHOR)

Revolutionizing Reinforcement Learning Framework for Diffusion Large Language Models

Sep.2025

<https://github.com/Gen-Verse/dllm-rl>

- **300+ stars on Github, 2k+ downloads on huggingface within 2 months**
- Proposed the first trajectory-aware reinforcement learning for diffusion language model, and a diffusion-based value model which can enhance training stability and naturally integrates process rewards.
- Employing this framework, we obtain 4B-scale diffusion language models that outperform 7B LLMs on reasoning tasks.

Co-Evolving LLM Coder and Unit Tester via Reinforcement Learning

May.2025

<https://github.com/Gen-Verse/CURE>

- **NeurIPS 2025 Spotlight**
- Proposed a general co-evolving reinforcement learning and apply it on coding agent setting with a simple yet effective statistics modeling.
- Using this framework to derive the state-of-the-art coding large language models.

Statistical Inference for Balance in Signed Networks

Oct.2024

<https://arxiv.org/pdf/2409.06172>

- **Biometrika major revision currently, main theorems developed by myself**
- Introduced a representation theory for signed networks, marking its inaugural proposal.
- Proposed a novel statistical inference method for assessing balance level in signed networks.

Multi-source Hierarchical Clustering for Electronic Health Records

May.2025

<https://arxiv.org/pdf/2509.06576>

- **JASA under review currently**

• Proposed a framework for automatically constructing a hierarchical knowledge graph for electronic health record (EHR) codes.

Mastering LLM Agent Workflows via Score-based Preference Optimization

Jan.2025

<https://arxiv.org/pdf/2502.04306>

- Proposed **the first** agentic workflow reinforcement learning framework.

MAIN HONORS

- *McCormick Fellowship from UChicago (first among 2024 first year PhD)* 2024
- *Chinese Mathematics Competitions, First Prize* 2022
- *National Scholarship (top1%, highest scholarship from Ministry of Education of China)* 2021
- *National High School Mathematics League, First Prize* 2019

TEACHING ASSISTANT

Introduction to Bayesian Data analysis (Winter 2025) Jan.2025 - Mar.2025
• Instructor: **Prof. Fei Liu** *Statistics Department, University of Chicago*
• 20 students enrolled.

Mathematical Analysis (Fall 2022) Sep.2022 - Jan.2023
• Instructor: **Prof. Yi Cheng** *School of Mathematical Sciences, USTC*
• 120 students enrolled.

TECHNICAL STRENGTH

Programming: Python, C/C++, R, MATLAB, Mathematica, LATEX, Stata, PLINK