Ruiwen Zhou

skyriver@sjtu.edu.cn https://skyriver-2000.github.io

Education

Shanghai Jiao Tong University Sep 2022 – Mar 2025
M.Eng. student in Computer Science (GPA: 3.83 / 4.00) *Shanghai, China*
Thesis: Design and Evaluation of LLM Complex Reasoning Methods and Agents

Shanghai Jiao Tong University Sep 2018 – Jun 2022
B.Eng. in Information Engineering (GPA: 90.56 / 100, Rank: 2 / 143) *Shanghai, China*
Thesis: Graph Neural Network-Based Tabular Data Prediction

Interest

My interest lies in building powerful language models and AI agents that can automatically complete complex real-world tasks and inspire creative ideas / designs for humans. To achieve this, my recent research works mainly focus on:

- Evaluation and analysis of LLMs.
- LLMs with retrieval augmentation.
- LLM reasoning, planning, rule-following, and agents.

Preprints / Under Review

Is Risk-Sensitive Reinforcement Learning Properly Resolved? arXiv preprint
R. Zhou, *M. Liu, K. Ren, X. Luo, W. Zhang, and D. Li*

Publications

RuleArena: A Benchmark for Rule-Guided Reasoning with LLMs in Real-World Scenarios ACL 2025
R. Zhou, *W. Hua, L. Pan, S. Cheng, X. Wu, E. Yu, and W. Wang*

AntiLeak-Bench: Anti-Leakage Benchmark for LLMs by Contamination-Free Samples ACL 2025
X. Wu, L. Pan, Y. Xie, ***R. Zhou****, Y. Ma, M. Du, R. Mao, S. Zhao, A. Luu, and W. Wang*

TRAD: Enhancing LLM Agents with Step-Wise Thought Retrieval and Aligned Decision SIGIR 2024
R. Zhou, *Y. Yang, M. Wen, Y. Wen, W. Wang, C. Xi, G. Xu, Y. Yu, and W. Zhang*

Learning Enhanced Representations for Tabular Data via Neighborhood Propagation NeurIPS 2022
K. Du, W. Zhang, ***R. Zhou****, Y. Wang, X. Zhao, J. Jin, Q. Gan, Z. Zhang, and D. Wipf*

Experience

NLP Group (UC Santa Barbara) Jul 2024 – Dec 2024
Visiting Student, Advised by: Prof. William Yang Wang *Santa Barbara, U.S.*

- Proposed a challenging benchmark (**RuleArena**) from real-world scenarios to evaluate LLMs' ability in rule-guided reasoning, and conducted extensive analysis to uncover systematic issues that limit LLM performances.
- Revealed that: 1) existing state-of-the-art LLMs, mostly fail on our complex rule-guided reasoning tasks; 2) LLMs struggle to integrate multiple rules or facts cohesively and are prone to distraction by irrelevant information; and 3) common failure modes include inadequate rule recall, improper usage of similar rules, and computation errors.
- Participated in the design and data collection of **AntiLeak-Bench**, which aims to prevent data contamination through automatically constructing benchmarks with updated real-world knowledge.
- **Two papers accepted at ACL 2025.**

APEX Lab (Shanghai Jiao Tong University) Jan 2021 – Mar 2025
Student Researcher, Advised by: Prof. Weinan Zhang *Shanghai, China*

- I worked as a student researcher under the supervision of Prof. Weinan Zhang.
- **I led the projects at MSRA and CPIC, and participated in the project at AWS when I worked in APEX Lab.**

China Pacific Insurance Company (CPIC) Feb 2023 – Feb 2024
Student Leader of a Collaboration Project *Shanghai, China*

- Revealed that existing trajectory-wise few-shot LLM agents suffer from plausible expert demonstrations due to retrieval with task meta-data and noise from many irrelevant steps in expert trajectories.
- Proposed a step-wise demonstration retrieval and prompting method (**TRAD**) to better solve sequential decision making tasks with LLMs, which achieves state-of-the-art performances on ALFWorld and Mind2Web benchmarks.
- **One paper accepted at SIGIR 2024.**

Amazon Web Service (AWS) Feb 2022 – Feb 2023
Research Intern, Advised by: Quan Gan *Shanghai, China*

- As existing retrieval-augmented tabular prediction models ignored either column-wise (across features) or row-wise (across samples) interaction, we aimed to develop a novel model architecture to unify both interactions and enhance the performance on various tabular prediction tasks.
- Participated in design and implementation of a novel tabular prediction model (**PET**) based on graph neural networks and relevant sample retrieval, which achieves state-of-the-art results on various tabular prediction benchmarks.
- **One paper accepted at NeurIPS 2022.**

Microsoft Research Asia (MSRA) Aug 2021 – Jan 2022
Research Intern, Advised by: Kan Ren *Shanghai, China*

- Revealed a common theoretical issue in existing distributional risk-sensitive RL algorithms - the absence of history return distributions in policy and value functions leads to optimization divergence.
- Proposed a history-dependent reinforcement learning algorithm (**Trajectory Q-Learning**), which achieves theoretical optimality and decent practical performance in risk-sensitive policy optimization under distortion risk measures.
- **One paper released on arXiv.**

Selected Awards

National Scholarship (Top 1 / 144)	2020
A-Class Excellence Scholarship (Top 1 / 144)	2020
B-Class Excellence Scholarship (Top 10%)	2019, 2021
Zhiyuan Honors Scholarship (Top 5%)	2019, 2020, 2021
First-Class Excellence Scholarship	2022, 2024
Huatai Securities Scholarship	2024

Talks

TRAD: Enhancing LLM Agents with Step-Wise Thought Retrieval and Aligned Decision Jul 2024
Oral Presentation at SIGIR 2024
Paper Talk at SIGIR 2024 AgentIR Workshop

Services

Organizer: Volunteer Host (SIGIR 2024 AgentIR Workshop)
Reviewer: ICML (2023), NeurIPS (2023), TPAMI

Skills

Programming: Proficient in Python, LaTeX; Capable of C, C++, Matlab
Languages: Proficient in Chinese, English (TOEFL: 106)