Xuanfei Ren

https://xuanfeiren.github.io/ xuanfeir@gmail.com github.com/xuanfeiren

Education

University of Wisconsin-Madison **Madison, USA**
PhD student, advised by Tengyang Xie, Department of Computer Sciences. *Aug. 2024 –*

University of Science and Technology of China **Hefei, China**
Bachelor of Science in Mathematics, specialization in Probability and Statistics. *Aug. 2020 – Jun. 2024*

Publication

- **Xuanfei Ren**, Allen Nie, Tengyang Xie, Ching-An Cheng.
  POLCA: Stochastic Generative Optimization with LLM
  *Submitted to **ICML 2026**.*
- **Xuanfei Ren**, Tianyuan Jin, Pan Xu.
  Optimal Batched Linear Bandits.
  *In Proc. of the 41st International Conference on Machine Learning (**ICML 2024**).*

Research Interests

My current research interests span broadly across **AI agents** and **reinforcement learning**, encompassing both theoretical foundations and experimental work.

Research Experience

Stochastic Generative Optimization with LLM

Work with Dr. Ching-An Cheng (Google Research), Dr. Allen Nie (Google Deepmind) and Prof. Tengyang Xie

- Formalized complex system tuning as a **stochastic generative optimization** problem, utilizing LLMs as optimizers to automate iterative refinement through numerical rewards and natural language feedback.
- Engineered a scalable framework featuring a **priority queue-based search strategy** to balance exploration-exploitation, systematically tracking solution histories under noisy evaluation conditions.
- Integrated an **ε-Net mechanism** to ensure parameter diversity and an **LLM Summarizer** for meta-learning, significantly reducing redundant trials and improving sample efficiency.
- Demonstrated state-of-the-art performance across diverse benchmarks: **τ-bench** (agent orchestration), **VeriBench** (code translation/Lean4), and **KernelBench** (high-performance CUDA kernels).
- Contributed to OpenTrace, an open-source library for end-to-end generative workflow optimization, enhancing its support for robust stochastic behaviors.

Offline Reinforcement Learning and Policy Evaluation Theory — University of Wisconsin-Madison

Advisor: Prof. Tengyang Xie *Sept. 2024 – Jan. 2025*

- Developed theoretical insights into value-based reinforcement learning (RL) algorithms.
- Theoretically analyzed estimation and approximation errors across various policy evaluation methods.

Optimal Batched Linear Bandits. — Duke University

Advisor: Prof. Pan Xu (Department of Computer Science, Duke University.) *Aug. 2023 – Jan. 2024*

- Devised an algorithm striving for asymptotic and non-asymptotic optimality simultaneously in the **linear bandits regret minimization** setting, an achievement previously unattained.
- Adapted the algorithm into a batched version with **provable least batch complexity**, extending applicablity to common real-world problems.
- Confirmed the algorithm's superiority over existing baseline methods through rigorous experimentation, showcasing its practical efficacy in linear bandits problems.

Community Service

Conference Reviewer: ICML, ICLR, RLC, ACL.

Skills

English
Programming: Python, PyTorch, R, LaTeX, MATLAB, Mathematica.