

# JIN PENG ZHOU

C: (607) 280-6520 | E: [jpzhou@cs.cornell.edu](mailto:jpzhou@cs.cornell.edu) | [jinpz.github.io](https://jinpz.github.io)

## EDUCATION

### Cornell University

**Ph.D. Candidate in Computer Science Supervised by Prof. Kilian Weinberger & Wen Sun** 2021/08-Present

- NSF NAIRR Pilot Award 2024/08
- NSERC CGS Doctoral Fellowship 2022/12
- Cornell University Fellowship 2021/08

### University of Toronto

**BASc in Engineering Science Machine Intelligence Option, Graduated with High Honours** 2016/09-2021/06

- Cumulative GPA: 3.95/4.0
- Undergraduate Thesis Supervisor: Prof. Roger Grosse
- Research Assistant under Prof. Nicolas Papernot and Prof. Scott Sanner

## SELECTED PUBLICATIONS ([Google Scholar](#))

- Jin Peng Zhou\***, Kaiwen Wang\*, Jonathan Chang, Zhaolin Gao, Nathan Kallus, Kilian Weinberger, Kianté Brantley, Wen Sun. Q#: Provably Optimal Distributional RL for LLM Post-Training. NeurIPS 2025
- Kaiwen Wang\*, **Jin Peng Zhou\***, Jonathan Chang, Zhaolin Gao, Nathan Kallus, Kianté Brantley, Wen Sun. Value-Guided Search for Efficient Chain-of-Thought Reasoning. NeurIPS 2025
- Jin Peng Zhou**, Seb Arnold, Nan Ding, Kilian Weinberger, Nan Hua, Fei Sha. Graders Should Cheat: Privileged Information Enables Expert-Level Automated Evaluations. EMNLP 2025
- Jin Peng Zhou\***, Christian Belardi\*, Ruihan Wu\*, Travis Zhang, Carla Gomes, Wen Sun, Kilian Weinberger. On Speeding Up Language Model Evaluation. ICLR 2025
- Jin Peng Zhou**, Charles Staats, Wenda Li, Christian Szegedy, Kilian Weinberger, Yuhuai Wu. Don't Trust: Verify - Grounding LLM Quantitative Reasoning with Autoformalization. ICLR 2024
- Jin Peng Zhou\***, Yuhuai Wu\*, Qiyang Li, Roger Grosse. REFACTOR: Learning to Extract Theorems from Proofs. ICLR 2024
- Zhenzhen Liu\*, **Jin Peng Zhou\***, Yufan Wang, Kilian Q. Weinberger. Unsupervised Out-of-Distribution Detection with Diffusion Inpainting. ICML 2023
- Albert Q. Jiang\*, Sean Welleck\*, **Jin Peng Zhou\***, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée Lacroix, Yuhuai Wu, Guillaume Lample. Draft, Sketch, and Prove: Guiding Formal Theorem Provers with Informal Proofs. ICLR 2023
- Ruihan Wu\*, **Jin Peng Zhou\***, Kilian Q Weinberger, Chuan Guo. Does Label Differential Privacy Prevent Label Inference Attacks? AISTATS 2023
- Jin Peng Zhou\***, Zhaoyue Cheng\*, Felipe Perez, Maksims Volkovs. TAFA: Two-headed Attention Fused Autoencoder for Context-Aware Recommendations. RecSys 2020

## EMPLOYMENT & TEACHING

**Research Scientist Intern, Meta MSL**, Menlo Park, United States 2025/05-2025/10

- Developed a fine-grained data filtering method for pre-training and mid-training datasets to improve LLM performance in mathematical reasoning tasks

**Student Researcher, Google DeepMind**, Mountain View, United States 2024/05-2024/12

- Developed a novel method to enhance the comprehensiveness and effectiveness of LLM evaluation, leading to the publication of an internal report and invited talk due to its strategic impact
- Collaborated with the Gemma Team to benchmark and evaluate the performance of Gemma 2 models

**Student Researcher, Google Research**, Mountain View, United States 2022/05-2023/05

- Contributed to the N2Formal team, focusing on the autoformalization of mathematical statements, definitions, and proofs
- Led a project utilizing theorem proving environments to ground LLM mathematical reasoning

**Machine Learning Researcher, TD Bank Layer6 AI**, Toronto, Canada 2019/09-2020/06

- Part of the team that won 2<sup>nd</sup> place in the Twitter RecSys 2020 challenge, leveraging a combination of feature engineering and deep language models to make predictions of user engagement

**Visiting Research Assistant (Mitacs Globalink Research Award)**, Supervisor: Prof. Vivienne Sze 2019/05-2019/09

**Energy-Efficient Multimedia Systems Group, Massachusetts Institute of Technology**

- Developed and improved efficient eye-tracking algorithms on embedded devices to track neurodegenerative diseases by measuring saccade latency and error rates. Deployed the application on iOS that serves many Alzheimer disease patients

**Instructor, Department of Computer Science, Cornell University** 2025/08-2025/12

- Taught CS1133: Introduction to Python Programming as the course instructor for 55 students