

JINMING NIAN

jmian@scu.edu ◊ (530)854-1122 ◊ Santa Clara, CA ◊ [Google Scholar](#) ◊ [Website](#)

SUMMARY

Third-year PhD candidate specializing in Information Retrieval, Recommendation Systems, NLP, and LLM Agents. Experienced in scientific programming with Python and PyTorch, with additional expertise in software development, cloud computing, and distributed systems. Actively publishing research in top-tier conferences on retrieval and LLM-driven methods.

EDUCATION

Santa Clara University

Ph.D. in Computer Science and Engineering

September 2023 – June 2028 (Expected)

Advisor: [Prof. Yi Fang](#)

M.S. in Computer Science and Engineering

September 2021 – June 2023

Relevant Courses: Machine Learning, Deep Learning, Information Retrieval, Natural Language Processing, Reinforcement Learning

University of California, Davis

September 2017 – June 2021

B.S. in Physics, Minor in Music

WORK EXPERIENCE

Data Scientist Intern, Amazon

June 2025 - September 2025

- Fine-tuned GPT-2 on sub-terabyte scale historical tabular shipment data to forecast Amazon fulfillment center shipment attributes (SLAM time, promise, and volume) during high-volume events such as Prime Day and Black Friday
- Improved SLAM-to-promise time prediction, reducing WAPE from 48.49% to 17.51% at a major Seattle fulfillment center
- Validated the approach's robustness and scalability across multiple fulfillment centers and temporal settings

Research Assistant, Santa Clara University

April 2023 - Present

- Recent work focuses on automatic prompt optimization for LLM agents and analysis of state-of-the-art deep research agent codebases; Previous projects are listed under **PUBLICATIONS**

Teaching Assistant, Santa Clara University

September 2025 - Present

- Information Retrieval (CSEN 169), Web Programming (CSEN 161)

PUBLICATIONS

• W-RAG: Weakly Supervised Dense Retrieval in RAG for Open-domain Question Answering

ACM SIGIR ICTIR 2025 – [paper](#), [code](#)

[Jinming Nian](#), Zhiyuan Peng, Qifan Wang, Yi Fang

• ELOQ: Resources for Enhancing LLM Detection of Out-of-Scope Questions

ACM SIGIR Resource & Reproducibility Track 2025 – [paper](#), [code](#), [dataset](#)

Zhiyuan Peng*, [Jinming Nian](#)*, Alexandre Evfimievski, Yi Fang

• Does Reasoning Introduce Bias? A Study of Social Bias Evaluation and Mitigation in LLM Reasoning

EMNLP 2025 (Findings) – [paper](#)

Xuyang Wu*, [Jinming Nian](#)*, Ting-Ruen Wei*, Zhiqiang Tao, Hsin-Tai Wu, Yi Fang

PROJECTS

DeepSeek R1-Zero for Reasoning Intensive Passage Re-ranking

February 2025 - Present

- Gathered training data: filtered hard queries where both lexical and dense retrieval methods fail; collected irrelevant documents that are consistently retrieved by strong retrievers as hard negative documents
- Redesigned reward modeling based on GRPO for the re-ranking task to encourage self-discovered CoT
- Leveraged Huggingface's Open-R1 as the base framework to train, test, and evaluate on the challenging BRIGHT dataset