

Zeinab Zolaktaf

✉ z.zolaktaf@gmail.com | 🏠 zainabzf.github.io/ | 🐙 [zainabzf](https://github.com/zainabzf) | 🌐 [zainabzolaktaf](https://zainabzolaktaf.com)

Summary

I build tools at the intersection of AI and software engineering, and lately I've been using agentic AI workflows to move much faster.

Skills: Applied Machine Learning & LLMs, Agentic AI Workflows & Tooling, LLM Evaluation & Fact-Checking, Generative AI Systems (Design to Production), Recommender Systems, NLP & Text Classification, ML Experimentation & Evaluation Frameworks, Cross-Team Technical Leadership.

Tools: Python, AWS, SQL.

Education

PhD in Computer Science, University of British Columbia

2012 - 2019

Vancouver, Canada

- Developed a novelty-aware recommendation framework for data exploration, countering the popularity bias of standard recommenders
- Built machine learning models that predict query properties (e.g., answer size, error type) before execution, to help users tune and optimize queries
- Developed statistical estimation techniques for analyzing large or conflicting query answers in data integration settings
- [Facilitating User Interaction With Data](#), PhD Thesis

MSc in Computer Science, Dalhousie University

2009 - 2012

Halifax, Canada

- Applied probabilistic modeling to community-based question-answering services, studying user expertise, participation, and content organization
- [Probabilistic Modeling in Community-Based Question Answering Services](#), MSc Thesis

BSc in Computer Engineering (Software Engineering), Isfahan University

2004 - 2008

Isfahan, Iran

- Built autonomous soccer-playing agents for the "UI-AI" RoboCup team over three years - my entry point into AI
- Sport Software Recommendation, BSc Thesis

Experience

Applied Scientist II

2022 - Present

AMAZON

Vancouver, Canada

- Built a governed, reusable agentic research workflow (spec-driven author/verifier loop with persistent project memory) that compounds across projects and transfers to other teams
- Used agentic workflows to ramp up in an unfamiliar domain and deliver an LLM judge for item analysis — approximating continuous variables via prompting
- Ideated, architected, and built an internal LLM experimentation and evaluation toolkit — proved 13x faster time-to-production, then partnered with engineering teams to scale it to 10 production generative AI applications across teams
- Built reusable, scalable reference-based LLM judges for fact-checking, used across many projects
- Led cross-team science and engineering efforts for an LLM-based recommender

Lead Applied Scientist

2021 - 2022

ELPHA SECURE

Vancouver, Canada

- Led cross-team collaborations to build ML products for cybersecurity, including classifiers for encrypted command detection, suspicious login detection, and URL phishing.
- Led the development of an ML pipeline to expedite ML application creation and experimentation, accelerating and streamlining how new ML products were built and deployed.
- Mentored 2 junior data scientists, providing guidance on ML model development and best practices.
- Automated ML model deployment with Kubernetes, featuring lightweight and scalable scripts.

Applied Data Scientist

2020 - 2021

GEORGIAN

Toronto, ON

- Performed an extensive error analysis of an existing ML product, helping business stakeholders and software engineers understand and analyze individual pipeline components.
- Designed and implemented NLP models for hierarchical text classification in the financial and marketing domain, including multi-label, multi-class, and translation models. Achieved a 188% improvement in F-score (aggregated across hierarchy).

Machine Learning Scientist

2019 - 2020

EHS&I

Vancouver, Canada

- Built NLP deep learning models for document intelligence in the environment, health, and safety domain.
- Designed data augmentation and synthetic data generation methods, greatly reducing manual labeling.
- Designed and developed an ML framework for deploying and evaluating pretrained language models for classification tasks; selected model achieved a 23% F-score improvement over the existing model.

Machine Learning Consultant

2021

OLYNS

San Francisco, USA

- Scientific verification of existing ML models and AWS deployment support.

Publications

- 2023 [Workload-Aware Query Recommendation Using Deep Learning](#), EDBT
- 2020 [Facilitating SQL Query Composition and Analysis](#), SIGMOD
- 2019 [Improvement of SQL Recommendation on Scientific Database](#), SSDBM
- 2018 [A Generic Top-N Recommendation Framework For Balancing Accuracy, Novelty, and Coverage](#), ICDE
- 2018 [Bridging The Gap Between User-centric and Offline Evaluation Of Recommendation Systems](#), UMAP
- 2017 [Facilitating User Interaction With Data](#), PhD@VLDB
- 2015 [Extracting Aggregate Answer Statistics for Integration](#), EDBT
- 2012 [Finding Expert Users in Community Question Answering](#), WWW
- 2011 [Modeling Community Question Answering Archives](#), NeurIPS