Qianfeng Wen

Toronto, Canada

647-894-7316 qianfeng.wen@mail.utoronto.ca LinkedIn Github Google Scholar

Education

University of Toronto **Sep.2026 –**
MSc in Computer Science, Supervisor: Ashton Anderson *Toronto, Canada*

University of Toronto **Sep. 2022 – Jun. 2026**
Bachelor. Computer Science Specialist, focus on AI *Toronto, Canada*
GPA 3.91/4.0

Publications (* denotes equal contribution)

SafeGEO: Understanding Generative Engine Optimization Risks in Recommendation Agents *2026*

Qianfeng Wen*, Yifan Simon Liu*, Xin Liu*, Difan Jiao, Blair Yang, Junda Wu, Zhenwei Tang
arXiv preprint; project page

Bayesian Active Learning with Gaussian Processes Guided by LLM Relevance Scoring for Dense Passage Retrieval *2026*

Junyoung Kim, Anton Korikov, Jiazhou Liang, Justin Cui, Yifan Simon Liu, ***Qianfeng Wen***, Mark Zhao, Scott Sanner
Findings of the Association for Computational Linguistics: ACL 2026

ThinkTwice: Jointly Optimizing Large Language Models for Reasoning and Self-Refinement *2026*

Difan Jiao*, ***Qianfeng Wen****, Blair Yang, Zhenwei Tang, Ashton Anderson
COLM 2026 under review

Grounded Chess Reasoning in Language Models via Master Distillation *2026*

Zhenwei Tang, ***Qianfeng Wen***, Seth Grief-Albert, Yahya Elgabra, Blair Yang, Honghua Dong, Ashton Anderson
COLM 2026 under review

A Simple but Effective Elaborative Query Reformulation Approach for Natural Language Recommendation *2025*

Qianfeng Wen*, Yifan Liu*, Justin Cui*, Joshua Zhang, Anton Korikov, George-Kirollos Saad, Scott Sanner
arXiv:2510.02656

Natural Language Recommendation via Multimodal Item Scoring Using Gaussian Process Regression with LLM Relevance Judgments *2025*

Yifan Liu*, ***Qianfeng Wen****, Jiazhou Liang*, Mark Zhao*, Justin Cui, Anton Korikov, Armin Toroghi, Junyoung Kim and Scott Sanner
Findings of the Association for Computational Linguistics: ACL 2026

ChessQA: Evaluating Large Language Models for Chess Understanding *2025*

Qianfeng Wen, Zhenwei Tang, Ashton Anderson
COLM 2026 Under review

MA-DPR: Manifold-aware Distance Metrics for Dense Passage Retrieval *2025*

Yifan Liu*, ***Qianfeng Wen****, Mark Zhao*, Jiazhou Liang, Scott Sanner
Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP Main)

A Comparative Study of Static and Contextual Embeddings for Analyzing Semantic Changes in Medieval Latin Charters *2024*

Yifan Liu, Gelila Tilahun, Xinxiang Gao, ***Qianfeng Wen*** and Michael Gervers
COLING'25 Workshop: LoResLM (oral)

Elaborative Subtopic Query Reformulation for Broad and Indirect Queries in Travel Destination Recommendation *2024*

Qianfeng Wen*, Yifan Liu*, Joshua Zhang, George Saad, Anton Korikov, Yury Sambale and Scott Sanner
RecSys '24 Workshop: ROEGEN

Monte-Carlo Tree Search for Behavior Planning in Autonomous Driving *2024*
Qianfeng Wen, Zhongyi Gong, Lifeng Zhou and Zhongshun Zhang
IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR) '24 (oral)

Working Experience

ZBot Technology **October 2025 – Present**
Co-founder & CTO *Toronto, Canada*

- Led development of a Export oriented AI agent system.

Air Canada **June 2024 – August 2024**
AI Engineer Intern (in collaboration with D3M Lab @ UofT) *Toronto, Canada*

- Led development of a travel recommender system that adapts to free-form user queries.
- Curated and released **TravelDest**, the first benchmark dataset for travel search, combining WikiVoyage content with insights from travel experts.
- Developed an LLM-based query reformulation + dense retrieval pipeline, improving R-precision by **12%** over baseline.

Bosch **May 2023 – August 2023**
Autonomous Driving Intern *Shanghai, China*

- Contributed to Bosch's *Wave3* autonomous driving program with a focus on motion planning algorithms.
- Designed an MCTS-based decision-making framework for complex urban/highway scenarios; achieved a **98%** success rate with planning latency $<$ **0.2**s on project benchmarks.

Research Experience

CSS Lab (University of Toronto) **May 2025 – June 2026**
Undergraduate Researcher; Supervisor: Ashton Anderson *Toronto, Canada*

- Created **ChessQA**, a benchmark to evaluate large language models on comprehensive chess understanding (tactics, strategy, notation, and explanation quality).
- Explored supervised fine-tuning and reinforcement learning to train and distill chess-specialized LLMs for reasoning and move justification.

D3M Lab (University of Toronto) **May 2024 – March 2026**
Summer Intern & Undergraduate Researcher; Supervisor: Scott Sanner *Toronto, Canada*

- Built Retrieval-Augmented Generation (RAG) QA pipelines combining sparse and dense retrieval (FAISS) with LLMs; developed LLM-driven query reformulation **EQR** and user-intent modeling to improve retrieval relevance.
- Proposed **Manifold Distance**, a training-agnostic metric inspired by Manifold Ranking, outperforming cosine/dot-product measures on Dense Passage Retrieval (DPR) benchmarks.
- Introduced **GPR-LLM**: a Gaussian Process Regression approach with LLM relevance judgments to model multimodal item relevance, consistently exceeding DR, cross-encoder, and pointwise LLM rerankers across four NLRec datasets and two backbone LLMs.

The Matter Lab (University of Toronto) **April 2024 – January 2026**
Undergraduate Researcher & Full-time Summer Intern; Supervisor: Alán Aspuru-Guzik *Toronto, Canada*

- Explored multi-agent debate (MAD), Chain-of-Thought (CoT), and Tree-of-Thought (ToT) methods to strengthen LLM logical reasoning; analyzed whether LLMs exhibit schema-based learning behaviors.
- Investigated compositional training methods (e.g., Ito density sampler) for diffusion models on 2D-ARC problems.

Projects

Decision-making in Autonomous Driving using Monte Carlo Tree Search (MCTS) | *Matlab* **90 Stars on github**

- Designed and implemented a comprehensive decision-making framework employing MCTS for diverse autonomous driving scenarios.
- Successfully validated the framework's effectiveness in varied conditions, ranging from complex urban intersections to highway exits.

Professional Service

Reviewer, IEEE International Conference on Robotics and Automation (ICRA) 2025

Honors & Awards

2026. Vector Scholarship in Artificial Intelligence. Vector Institute, Toronto. Research. $17,500

2025. Department of Computer Science Research Award. Department of Computer Science, University of Toronto. Institutional. Research. $8,000.

2023-2025. Dean's List Scholar. University of Toronto. Institutional. Academic. $0.

2024-2025. Chancellor's Scholarship. Trinity College, University of Toronto. Institutional. Academic. $500.

2024. Educational Training Stipend. Department of Mechanical & Industrial Engineering, University of Toronto. Institutional. Research. $2,500.

Technical Skills

Languages: Python, Java, C, Matlab, R
Technologies/Frameworks: verl, vLLM, SGLang, Prompt Engineering (OpenAI, LangChain), Pytorch, SkLearn, GitHub, Matlab Tool Boxes