Amir David

amirdavid.ca | Google Scholar | GitHub | LinkedIn
Waterloo, ON, Canada | Open to Toronto, Montreal, Remote Canada

SUMMARY

ML engineer and UWaterloo MMath graduate with **4+ years** of applied ML research and development experience across industry and academia. Built model training/evaluation pipelines and ML-backed services spanning computer vision, LLM evaluation, and RAG systems. Research background in privacy-preserving ML and LLM evaluation.

WORK EXPERIENCE

University of Waterloo - ML Research Assistant **01/2024 - 01/2026**

- Developed data augmentation methods for privacy-preserving ML model evaluation, enabling organizations to benchmark external model providers on proprietary datasets without sensitive data leakage. Collaborator: **RBC**.

Canadian Institute for Cybersecurity - ML Engineer **09/2021 - 12/2023**

- Built and deployed a privacy-preserving split-learning pipeline in PyTorch for secure computer-vision inference, hardened against attribute-inference and reconstruction attacks. Collaborator: **Huawei Canada**.
- Owned the training, evaluation, and inference pipeline for a YOLOv5 computer-vision model used in privacy/security research, including data preprocessing, tuning, and packaging; achieved 90% mAP. Collaborator: **Huawei Canada**.
- Built Python/Django REST services for a microservices-based security application that integrated ML inference into analyst workflows. Collaborator: **Rogers Communications**.

PROJECTS

Technical: Built a document Q&A agent with FastAPI, ChromaDB, and Sentence Transformers, supporting OpenAI and local GPU-backed inference (llama-cpp), hybrid semantic + keyword retrieval, source-cited answers, and faithfulness checks via NLI entailment, citation coverage, and claim verification. GitHub: https://github.com/adavid13/ai-agent.

Thesis: Showed that AUROC-based evaluation overstates LLM-text detector utility at deployment-relevant low-FPR thresholds, quantifying cross-generation transfer failure across 10 OpenAI and Anthropic model generations and 3 domains. Built a two-stage DeBERTa pipeline with hard example mining and Neyman-Pearson calibrated thresholding that reduced document-level FPR from ~1% to 0.18% while maintaining 90% TPR.

SKILLS

Languages: Python, SQL, JavaScript, Java, C, Bash
ML: PyTorch, LLMs, NumPy, Pandas, Transformers, RAG, Deep Learning, Computer Vision, Model Evaluation, HuggingFace, Fine-tuning, W&B, Model Deployment, Prompt Engineering, Embeddings
Backend & Dev: FastAPI, Django, Flask, REST APIs, Docker, SDLC, Git, GitLab CI/CD, Linux, AWS

EDUCATION

Master of Mathematics in Computer Science (GPA: 94%) **01/2024 - 01/2026**
University of Waterloo, Waterloo, Canada
Supervisor: Professor Florian Kerschbaum
Recognition: Graduate Excellence Award in CS (2024), Math Domestic Graduate Student Award (2024)

Bachelor of Science in Software Engineering **09/2016 - 04/2021**
University of New Brunswick, Fredericton, Canada
Recognition: Dean's List for the Faculties of CS and Engineering (2018-2021)

RESEARCH

- Peer reviewer - Forty-Third International Conference on Machine Learning (**ICML 2026 - Gold Reviewer**)
- Personally Identifiable Information Detection (**Knowledge-Based Systems, Oct 2025)**
- Towards Privacy-Preserving Split Learning (**Internet of Things, May 2025**)
- A Survey on Deep Learning in Edge-Cloud Collaboration (**Knowledge-Based Systems, Feb 2025**)