

Summary

Machine Learning Engineer / Senior Software Engineer with 5+ years building **ML platforms, RAG systems, and cloud-native ML software** across finance, healthcare, and enterprise environments. Strong background in backend systems, distributed services, and applied AI/ML. Currently pursuing a Master of Computer Science at UT Austin focused on AI/ML systems. Open to roles in **Machine Learning Engineering, ML Platform/MLOps, AI Engineering, or Backend/Infrastructure Engineering**.

Technical Skills

- **Languages:** Python, Go, SQL, TypeScript/JavaScript
- **ML/LLM:** PyTorch, HuggingFace, scikit-learn, MLflow, LangChain, LangGraph, Langfuse, RAG
- **Backend:** FastAPI, gRPC, REST, async systems, microservices
- **Data:** PostgreSQL, SQL Server, Snowflake, Redis, MongoDB, Milvus, DuckDB
- **Cloud/Infra:** AWS, Azure, GCP, Docker, Kubernetes, Terraform, GitHub Actions, Jenkins, Kafka
- **Security/Observability:** IAM, OAuth2, Vault, OPA/Sentinel, Prometheus, Grafana, Splunk

Research Experience

- **Dell Medical School (NPRQI) - University of Texas at Austin** Austin, TX
Research Engineer (Applied ML & Data Systems) Jan 2026 – Present
- Building **Python/SQL pipelines** over large-scale **Ascension Health** pediatric emergency care data to support **OMOP-based analytics** and measure development for clinical research workflows.
Walker Department of Mechanical Engineering (GAMMA Lab) - University of Texas at Austin Austin, TX
Graduate Researcher (LLM Systems) Jan 2026 – Present
- Researching **SFT** and **RFT** for **LLM-based scientific agents**, building infrastructure for experiment state machines, evaluation, and controller experimentation in **X-ray diffractometer** workflows.

Professional Experience

- **BMO Capital Markets (BMO Harris US)** Toronto, Canada
Application Software Engineer, AI/LLM Platform (Contract) Sep 2025 – Apr 2026
- Architected an **AI-assisted enterprise data governance platform** for Capital Markets, consolidating data profiling, classification, SQL translation, and data quality functions into a shared self-service system, reducing manual effort.
- Engineered a central **RAG knowledge base** for enterprise data and metadata use cases, implementing graph-based orchestration, prompt tracing, and evaluation over documentation and internal data assets.
- Built **AI-enhanced tools** for metadata reasoning, natural-language querying, and documentation generation, translating structured frontend inputs into backend **Langfuse**-tracked LLM requests over enterprise data sources.
- **Xanthus Security** Toronto, Canada
Software Engineering Consultant Jun 2025 – Aug 2025
- Built a **client-facing OSINT platform** integrating Shodan/IntelX with an **LLM-powered RAG pipeline** on **AWS Bedrock**, delivering PlexTrac-ready reports for real-world security incident workflows.
- **AI Spectral** Toronto, Canada
Machine Learning Engineer (Contract) Jul 2024 – May 2025
- Deployed **medical computer vision models** to **500+ Jetson Orin Nano** edge devices, building secure OTA delivery and offline-safe inference services using **Python, Docker, and Unix domain sockets** for production edge deployments.
- Designed distributed edge orchestration and ML training pipelines (**PyTorch Lightning, AzureML, MLflow**), improving reliability and p95 latency by **~50%** in resource-constrained edge environments.
- Owned observability for a distributed edge ML fleet using **Prometheus/Loki/Grafana**, cutting response time by **~40%**.
- **CIBC** Toronto, Canada
Senior Software Engineer (Contract) Nov 2024 – Apr 2025
- Developed a secure internal **LLM assistant** and scaled adoption from **5K** to **35K+** employees by integrating persistent chat history and **RAG** over Risk, Legal, and Audit data with RBAC and PII redaction for enterprise knowledge access.
- Created a **Teams AI meeting agent** POC with **Azure OpenAI** and Azure AI Speech custom lexicon, generating summaries and action items from meetings, reducing manual note-taking by **95%** for internal users.
- **Royal Bank of Canada (City National Bank)** Toronto, Canada
Senior Software Engineer May 2023 – May 2024
- Built a **multi-cloud** security exemption platform and automated creation, renewal, and expiry workflows using cloud-native SDKs/APIs across **Azure/AWS/GCP**, standardizing evidence capture for yearly federal audits.
- Shipped an **AI-driven change monitor** that ingested cloud/vendor advisories and mapped them to impacted controls and exemptions, cutting manual review for security architects by **dozens of hours per year**.
- Delivered dashboards and alerts for expiries and drift, enforced approval rules with **Azure Policy** and **HashiCorp Sentinel**.
- **Just Web Agency** Toronto, Canada
Software Engineer Aug 2021 – Feb 2023
- Built client-facing React.js dashboards and **scraping + lead-ingestion services** for a B2B lead-generation platform, driving a **3x increase in qualified leads**.

Education

- **University of Texas at Austin** Austin, Texas
Master of Computer Science Sep 2025 – Dec 2026 (Expected)
- **Carleton University** Ottawa, Canada
Bachelor of Computer Science (Honours), Minor in Business Sep 2019 – Dec 2022
- Graduated with Distinction, Dean's List