

# Umad Akram

[akramumad@gmail.com](mailto:akramumad@gmail.com) | [linkedin.com/umad](https://www.linkedin.com/umad) | [github.com/umad](https://github.com/umad) | [umadakram.com](https://umadakram.com)

## EDUCATION

---

### Ontario Tech University

Bachelor of Science in Computer Science (Co-op)

Oshawa, ON

Expected Graduation: April 2028

## EXPERIENCE

---

### Shopify

Software Engineering Intern

- Incoming Summer 2026.

May 2026 – Aug 2026

Toronto, ON

### Creation

Software Engineer

- Developed a **Python** backend deployed on AWS that translated natural-language prompts into parametric CAD workflows, generating production-ready models and attracting a waitlist of 300+ engineers.
- Built an NLP pipeline using Claude API to parse unstructured design prompts and generate manufacturable CAD models through the SolidWorks API, supporting **100+** automated design generations.
- Engineered a RAG pipeline with **1,000+** PostgreSQL records to improve model accuracy, consistency, and iteration speed by retrieving past successful responses.

Oct 2025 – Present

Toronto, ON

### Canadian Imperial Bank of Commerce

Software Engineering Intern - Co-op

- Automated Kubernetes infrastructure provisioning using **Go** and Terraform across 150+ microservices, reducing environment setup time from **5 days to under 30 minutes**.
- Designed a scalable LLM serving layer using FastAPI, Redis caching, and **Kubernetes** autoscaling to handle **25K+** daily prompt requests, reducing average response latency by 40ms.
- Implemented prompt logging, schema validation, and evaluation dashboards using OpenTelemetry and PostgreSQL, enabling structured analysis of hallucination patterns across **150K+** weekly LLM responses.

Jan 2025 – Aug 2025

Toronto, ON

### Canadian Imperial Bank of Commerce

Software Engineering Intern

- Built a **Python** based log parsing framework to process **1M+** weekly network events, enabling structured analysis of large-scale telemetry and faster root-cause identification during incidents.
- Implemented **CI/CD** pipelines using Azure DevOps across 20+ services, configuring build triggers, environment specific release stages, and deployment approvals to support reliable production releases.
- Resolved a critical bug in a Python script, averting potential damages of over **\$10 million** for the bank.

Sept 2024 – Dec 2024

Toronto, ON

## PROJECTS

---

### TradeGuard | Python, Typescript, React, Gemini API, MongoDB

- Engineered a real-time data ingestion pipeline in **TypeScript** using the Kalshi API to process **100,000+** active markets, storing market state, trade history, and derived features in MongoDB.
- Built an ML anomaly detection system using 14 quantitative signals, with a **Python** inference pipeline and Google Gemini API to flag suspicious behavior and generate a 0–100 risk score.

### DocFlow | Hack the North (Top 10 YC Finalist) | Python, Node.js, DynamoDB, Gemini API

- Engineered a documentation automation tool with Node.js, GitHub Webhooks, and Gemini API that converts PR merges into production-ready docs, reducing documentation time from **2 hours to 10 minutes per PR**.
- Architected an automated code review pipeline with **Python** and DynamoDB that flagged 100+ bugs, security vulnerabilities, and logic flaws across **50+** merges.

### Ride Vision | Go, TypeScript, React.js, Node.js, PostgreSQL, Docker, Nginx

- Developed a full-stack ride-hailing simulator with Go, Node.js simulation engine, and React Frontend, supporting real-time driver–rider matching, routing, and surge pricing for **1,000+** simulated trips.

## TECHNICAL SKILLS

---

**Languages:** Python, Go, JavaScript/TypeScript, SQL, C++

**Frameworks/Libraries:** React.js, Node.js, Flask, PyTorch

**Tools/Technologies:** Git, AWS, Azure, Kubernetes, Terraform, Docker, Redis, PostgreSQL, DynamoDB, MongoDB