

Joseph Barbosa

josephbarbosa.com | 647-328-9032 | josephbarbosa416@gmail.com | linkedin.com/in/joseph-c-barbosa | github.com/joeyb007

EDUCATION

University of Waterloo & Wilfrid Laurier University

GPA: 3.98 / 4.0 (92%)

BSc Computer Science (UW) & BBA Business Administration (WLU)

Sep. 2025 – Present

Relevant Coursework: Data Structures and Algorithms, Linear Algebra, Calculus I-II, Optimization, OOP, Probability

EXPERIENCE

Software Engineering Intern

May 2026 – Aug. 2026

Scotiabank

Toronto, ON

- Full-stack engineer on the CBE production banking platform (React, Spring Boot, Docker); leading an agentic AI QA initiative by building an LLM agent that analyzes code diffs and generates targeted regression tests in the CI/CD pipeline

Software Engineering Intern

Jan. 2026 – Apr. 2026

ESGTree

Waterloo, ON

- Built a production NL-to-SQL agent over a **98-table** MySQL DB using a custom ReAct loop on the Anthropic Messages API with FAISS schema retrieval and FK-aware join hints; designed an eval suite of **77** user queries from production logs achieving **90.9%** accuracy, with ablations showing self-correction and retrieval contributed **+13.0pp** and **+7.8pp** respectively
- Engineered backend with read-only DB privileges, regex SQL validation, and prompt-level constraints; enforced multi-tenant isolation via JWT; reduced post-query latency by delegating narrative summarization and chart spec generation to Haiku
- Extended the agent with Docker sandboxed Python execution (Opus for code generation, pandas/scipy/sklearn pre-loaded), enabling trend projection and statistical analysis beyond SQL; architected a 4-model cost/latency tier: Sonnet for reasoning, Opus for code generation, Haiku for planning, summarization, and chart spec, matching model capability to task complexity

Technical Project Manager

Sep. 2025 – Present

LuminWell

Waterloo, ON

- Led an **8+ member** team of graduate students and undergraduate interns to build and deploy a Gemini-powered peer-matching platform (MERN stack), securing **\$2.5K** in external funding for mental health access initiatives

Software Development Intern

Jun. 2024 – Dec. 2024

FuturIQ

Brampton, ON

- Built a production ML-powered real estate platform integrating a full-stack web application (Flask, PostgreSQL, JavaScript) with an automated price prediction pipeline that scraped **1000+** live housing market data points via Selenium, trained models on dynamically collected market subsets, and served predictions through the application backend
- Implemented an ML training and evaluation pipeline using train-test split, cross-validation, and RMSE, comparing linear and polynomial regression models and deploying a regularized polynomial model to improve generalization

PROJECTS

Scholar – Agentic AI Research Assistant. tryscholar.com

Python, FastAPI, asyncio, OpenAI API, Pydantic, OpenAlex, MCP, Next.js, NextAuth.js, Postgres

- Architected a recursive retrieval pipeline over **200M+ papers** (OpenAlex) with parallel asyncio fan-out, depth-limited search, and cost-bounded API ceilings; shipped to **300+ users** at tryscholar.com
- Engineered hallucination-resistant synthesis with post-generation citation validation, 3-retry query replanning, and LLM-inferred follow-up detection with full session context
- Delivered as CLI, **MCP server** (Claude Desktop / Cursor), and full-stack web app (FastAPI SSE, Next.js with live token streaming, Postgres persistence, Google OAuth); **55 unit tests** with mocked LLM I/O

Studeal – Agentic Deal Discovery System

Python, FastAPI, LangGraph, Next.js, PostgreSQL + pgvector, Redis, Celery

- Designed a multi-stage, LangGraph-based agentic retrieval system for deal discovery (retrieval, extraction, scoring)
- Built an asynchronous pipeline, parallelizing search, scraping, and inference, including semantic deduplication (pgvector)
- Implemented a scoring agent with market context for daily personalized deal matching from natural language watchlists
- Achieved **90%+ link resolution accuracy** by reconstructing missing product URLs using LLM-guided extraction
- Shipped production system with rate-limited APIs, Celery scheduling, and Stripe access control, plus Next.js UI

PantryPal – Fine-Tuned LLM for Constraint-Aware Generation

Python, FastAPI, Llama 3.2, QLoRA, DPO, XGBoost

- Fine-tuned Llama 3.2 (3B) with QLoRA on **10K instruction samples** for constraint-aware recipe generation
- Built the synthetic dataset using **18 XGBoost classifiers** over ingredient features, enabling DPO without human labels

TECHNICAL SKILLS

Languages: Python, SQL (PostgreSQL, SQLite), TypeScript, HTML/CSS, Java, C

AI / ML: LangChain, LangGraph, Hugging Face, XGBoost, OpenAI/Anthropic/Groq APIs, RAG, pgvector, ChromaDB

Frameworks: React, FastAPI, Flask, Next.js, Node.js, Celery

Libraries: pandas, NumPy, scikit-learn, Pydantic, SQLAlchemy, Selenium, BeautifulSoup, Matplotlib

Tools: Git, Docker, Redis, Bash, Figma, Postman, Supabase, Heroku, Railway, Claude Code

Certifications: Stanford Machine Learning Specialization (Supervised, Unsupervised, Reinforcement Learning)

Awards: President's Gold Scholarship, Loblaw National Scholarship, National Meritorious Award (IMMC), \$3,000+ in Case Competition Winnings (5+ finalist placements)