

ARIYAN AZAMI

Toronto, ON | +1 (437) 599-7338 | ariyan.azami84@gmail.com
ariyanazami.com | linkedin.com/in/ariyan-azami | github.com/AriyanAzami

EDUCATION

YORK UNIVERSITY

B.Eng., Computer Engineering • Toronto, ON

Sep 2025 – Apr 2029

EXPERIENCE

AUGOVA

Co-Founder & Co-President • Richmond Hill, ON

AI automation for small and mid-sized businesses. Voice agents, document intelligence, process automation, custom AI

Aug 2026 – Present

- Lead business development, translating each business's workflow into a tailored AI solution; oversee service agreements, service-level commitments, and the proof period for new clients.

APPLI SOLUTIONS

Software Developer • Toronto, ON

AI hiring and professional-marketplace platform. Python, Flask, Firestore, Google Cloud Run, Vertex AI, Neo4j, Next.js

Jan 2026 – Present

- Cut p50 endpoint latency **5.4x** by replacing an N+1 read pattern with fixed three-round-trip batch hydration, making response time flat in page size rather than linear.
- Designed and shipped a community-recommendation engine as **sole developer** across four production services, reaching **99.2% hit@1** and **0.996 MRR** against a **12.5% random baseline** on a synthetic 240-user, 48-community cohort.
- Measured realized score contribution across the engine's seven weighted signals and found **45% of the declared weight budget inert**, with embedding similarity dominating the ranking while the heaviest-weighted signal contributed only marginally.
- Implemented the platform's connections and followers social graph across **14 REST endpoints** with transactional duplicate prevention, server-side count aggregation, and shared cursor pagination.
- Shipped the Activity Score as a nightly Cloud Run job applying half-life exponential decay across dozens of weighted event types.

LABORATORY OF ADVANCED BIOTECHNOLOGIES, YORK UNIVERSITY

Machine Learning Research Assistant • Toronto, ON

Mar 2026 – Jul 2026

- Developed and trained U-Net segmentation models in PyTorch and TensorFlow to localize early-stage disease markers in medical imaging, tuning architecture depth and augmentation strategy for small clinical datasets.
- Wrote Python preprocessing pipelines normalizing heterogeneous imaging inputs, later extending them with Stable Diffusion synthesis.

RESEARCH

REAL-TIME OBSTACLE DETECTION DATASET

Second author • submitted for publication 2026

- Trained and benchmarked **all six nano-scale detectors from YOLOv8n through YOLO26n** on a 24,326-image / 40,195-box / 25-class benchmark under one identical protocol, reaching **0.889 recall, 0.925 precision, 0.882 mAP@0.50, 0.750 mAP@0.50:0.95, 1.98M parameters, and 79.7 FPS**.
- Established the paper's central result that **recall and mAP rankings disagree**, proving that selecting a detector on aggregate mAP does not yield the one missing the fewest obstacles.
- Ran field capture and annotation across several of the 13 dataset campaigns, contributing **255 of 255 released boxes** for Fire hydrant and regrowing Bench from 260 to 774 instances (**+197%**) after it failed the gate at **0.769 precision**.
- Contributed to a dataset now at **83,000+ Hugging Face** and **3,800+ Kaggle downloads** that reconciles **29 public collections onto one class index** alongside original street-level capture in Toronto and Tehran.

PROJECTS

PROSECUTO

Python • LangGraph • vLLM • Chroma • Next.js

- Owned the local-inference path at **NVIDIA Spark Hack Toronto**, routing generation to on-device llama-server and vLLM NVFP4 on GB10 hardware rather than a hosted API.

UNDERCUT

Next.js • TypeScript • FastAPI • Go • PostgreSQL

- Created the visualization layer for a used-car deal finder flagging underpriced GTA listings, charting total cost of ownership and price history against listed price.

SYNCAAGENT

Python • zero dependencies

- Built a multi-agent orchestrator running Claude Code, Gemini CLI, and Codex CLI on one task through markdown files on disk.

TECHNICAL SKILLS

Languages: Python, C++, TypeScript / JavaScript, Java, SQL, MATLAB

ML & Computer Vision: PyTorch, TensorFlow, Ultralytics YOLO, U-Net, Stable Diffusion, large language models (LLMs), LangGraph, vLLM, Vertex AI embeddings, Hugging Face, OpenCV, Roboflow, SAM 3, CVAT, multi-GPU / distributed training (PyTorch DDP)

Backend & Cloud: Flask, REST API design, Google Cloud Run / Tasks / Scheduler, Firestore, Neo4j, PostgreSQL, Redis, Docker, GitHub Actions, pytest

Frontend: Next.js, React, Tailwind CSS, data visualization

Embedded & Hardware: Arduino, C++ firmware, ArduPilot, PID tuning, CAD (SolidWorks, Fusion 360), Altium Designer PCB, sensor integration

AWARDS

ATF Robotics Cup, 1st Place National (5,000+ participants) – led software and systems engineering, delivering full C++ Arduino firmware and custom Altium PCBs.

FIRA RoboWorldCup, 2nd Place National and World Finals qualifier – served as technical operations manager, troubleshooting firmware live under competition conditions.