

Mohammed Naqi

mo.naqi321@gmail.com | linkedin.com/in/mohammed-naqi-2b98002b6 | github.com/Mo-Naq1

EDUCATION

University of Waterloo

Bachelor of Software Engineering | 85% Average

Expected 2029

Waterloo, ON

Relevant Coursework: Statistics, Linear Algebra, Data Structures & Algorithms, Logic and Computation

EXPERIENCE

Software Engineer

Boost Collective

May 2026 – Aug 2026

Toronto, ON

- Shipped fulfillment + payment features for a **67k**-artist platform processing **\$11M+** across 85k orders.
- Built the Stream Tracking system on a **Postgres** trigger across 10.6k campaigns and **26.5M** tracked streams.
- Pitched and built a buy-one-get-one free-song promo end-to-end that ran **\$182k** across **988** orders in 3 weeks.
- Kept unfilled bonus orders out of both admin queues via a **PostgREST** computed-column SQL function.
- Fixed the root cause that broke site-wide SSR (FCP **3.1s** → **1.4s**) and lifted a checkout from PSI **12** → **41**.
- Rolled out deny-by-default **RLS** across ~51 production tables, verified by **287** pgTAP privilege tests.

Machine Learning Engineer

WAT.ai (WorldFold) | [Source](#)

May 2026 – Present

Waterloo, ON

- Built a cloth-folding RL env in **MuJoCo** + **Gymnasium** with a 141-dim observation (cloth, robot, task).
- Shaped a dense corner-to-goal distance reward plus a sparse **+10** success bonus at fold score ≥ 0.85 .
- Defined a **14-dim** continuous action space for two 5-DOF SO101 arms, actuated via Jacobian least-squares IK.
- Trained **PPO** (Stable-Baselines3 MlpPolicy) on a flattened Dict→Box observation for **200k** steps.
- Faked flat-cloth grasping via a runtime **MuJoCo** WELD constraint and fixed a NaN contact blow-up on touch.

Software Engineer, Exoskeleton Project

UW BioTron Design Team | [Source](#)

Sep 2025 – Jan 2026

Waterloo, ON

- Built a **NiceGUI** + Plotly dashboard streaming **50 Hz** motor-current and joint-angle telemetry from an ESP32.
- Wrote a **Pandas** pipeline that flattens nested JSON telemetry into analysis-ready CSVs.

TECHNICAL SKILLS

Languages: Python, TypeScript, SQL (PostgreSQL), C/C++, Ruby, HTML/CSS

ML & Data: PyTorch, NumPy, Pandas, Scikit-learn, Gymnasium, MuJoCo, Stable-Baselines3, Hugging Face

Frameworks: Next.js, React 19, FastAPI, Tailwind CSS

Tools & Cloud: Docker, Supabase, Git, Vercel, TensorBoard

PROJECTS

Ocular Disease Classification (CNN) | [Source](#)

- Building a **PyTorch** computer-vision pipeline that classifies eye images into **6** categories (5 diseases + healthy).
- Training a multi-class **CNN** end-to-end for supervised disease classification across all 6 categories.
- Using **CrossEntropyLoss** with stratified train/validation splits to keep all 6 classes balanced during training.
- Evaluating with confusion matrices and per-class precision/recall to surface minority-class misclassifications.

PocketProf (AI Voice Learning Platform) | [Source](#) | [Live](#)

- Built a full-stack learning app on **React 19** + **FastAPI** delivering live, two-way voice tutoring sessions.
- Streamed PCM audio frames over **WebSocket** into **Smallest.ai Pulse** for live speech-to-text captions.
- Used **Gemini 2.5-flash** to turn raw lecture transcripts into clean, structured study notes with headings.
- Used **Gemini Vision** to match each spoken student question to whichever slide is currently on screen.

ESL Object Learner | [Source](#) | [Live](#)

- Built a fully client-side ESL vocabulary app in **vanilla JS/HTML/CSS** with zero backend or build step.
- Ran real-time object detection fully in-browser with **TensorFlow.js** + **COCO-SSD**, with no server calls.
- Wrote custom **IoU** and center-distance logic to dedupe overlapping detection boxes into single labels.
- Used scene-context size ratios to relabel visually similar objects the detector kept confusing.
- Scored pronunciation with the **Web Speech API** and **Levenshtein distance** against the target word.
- Built syllable-by-syllable TTS playback at adjustable speeds using the browser **SpeechSynthesis API**.