

Ryan Neil Alumkal

ryan.alumkal@mail.utoronto.ca [linkedin.com/in/ryanalumkal](https://www.linkedin.com/in/ryanalumkal) github.com/ryanalumkal

Education

University of Toronto

Expected May 2028

Bachelor of Applied Science (BASc), Electrical Engineering

Toronto, ON

- **Coursework:** Introduction to Machine Learning, Probabilistic Reasoning, Applied Fundamentals of Deep Learning, Computer Architecture (A-), Computer Organization (A), Computer Hardware

Skills

Languages: Python, C/C++, Verilog, SystemVerilog

Libraries: Pandas, NumPy, SciPy, Matplotlib, OpenCV, Transformers

Frameworks: PyTorch, Tensorflow, Scikit-learn

Other: Linux/Unix, Git

Experience

ASIC Team Associate

Oct 2025 - Present

IEEE University of Toronto Student Branch

Toronto, ON

- Mastered the full RTL-to-GDSII digital IC flow using the SKY130 open-source toolchain, including LibreLane, Open Circuit Design, and TinyTapeout.
- Executed end-to-end ASIC development, encompassing RTL design, functional simulation, logic synthesis, and physical layout for tapeout-ready designs.

Machine Learning Developer - Breast Cancer Detection

July 2025 - Jan 2026

University of Toronto Machine Intelligence Student Team (UTMIST)

Toronto, ON

- Systematically evaluated CNN, ResNet, and Vision Transformer architectures on 15+ medical imaging datasets; utilized hyperparameter optimization to significantly improve F1-scores for early-stage breast cancer detection.
- Engineered a reproducible research framework featuring headless tuning scripts and automated logging, ensuring full experimental traceability and zero data leakage across patient-level splits.

Projects and Coursework

Anime Portrait Generation Using GANs — *Python, PyTorch, GANs, CNNs, Git*

- Implemented a CycleGAN/Deep Convolutional GAN in PyTorch for artistic style transfer, training on a 20,000+ image dataset to convert real-world photos into anime-style images.
- Conducted systematic training and qualitative testing to optimize visual fidelity, achieving a 95.8% accuracy and an FID score of < 1000 for high-quality, realistic style transfer results.

Voice Controlled Rover — *Embedded Systems, C, RISC-V, DSP*

- Built a voice-activated rover on a RISC-V soft processor by digitizing 2-second microphone inputs with an ADC to process commands for movement and braking.
- Developed embedded C firmware with 10+ MATLAB-generated digital filters for real-time signal processing, sound classification, and motor control using GPIO, achieving sub-100ms actuation latency.

Car Object Detection — *Faster R-CNNs, PyTorch, Python, OpenCV*

- Created and fine-tuned a Faster R-CNN model in PyTorch to detect and localize car outlines on a labeled video dataset of 500+ clips (20,000+ frames), achieving a mean Intersection over Union (IoU) of 0.87.
- Preprocessed and augmented video frames, managed end-to-end training pipelines, and evaluated performance using IoU and Mean Squared Error (MSE), reducing bounding box localization error by 15%.

Data Cache Prefetcher Design with SimpleScalar — *Computer Architecture, C/C++*

- Implemented and evaluated L1 data-cache next-line, stride, and custom open-ended prefetchers in the SimpleScalar sim-cache framework, validating correctness with targeted microbenchmarks and AMAT calculations on compress/gcc/go workloads.
- Designed and integrated a GS + GHB-based data prefetcher that reduced the average L1 data-cache miss rate from to 2.06% while keeping hardware overhead modest.