

# Austin Luu

✉ austinowenluu@gmail.com

🏠 austinluu.me

🐙 austinluu

🌐 austin-luu

## Experience

---

### ViTAA Medical Solutions | SOFTWARE DEVELOPMENT ENGINEER II

Montreal, Canada | June 2024 — Present

- Designed and built a scalable machine learning platform with datasets in **AWS S3** and distributed training and inference on **EC2** HPC nodes, cutting new-model deployment time by 80% (5 days to 1 day).
- Architected a subagent-driven AI development pipeline integrating **MCP** servers with Confluence/Jira, coordinating 11 specialized agents across the SDLC and reducing development time by 88%.
- Built a model- and provider-agnostic dynamic context system that assembles, scopes, and prunes **LLM** context at runtime, cutting token consumption by 82–87% across 4+ projects.
- Established production monitoring, alerting, and automated recovery across all backend services with **CloudWatch**, raising coverage from 0 to 100% and cutting incidents by 80%.
- Owned architectural redesign of **Slurm**-managed data pipelines on **AWS EC2** with **Python** parallel execution, reducing runtime by 39%; an event-driven **AWS Lambda** workflow cut downtime by a projected 67%.
- Built reusable internal libraries standardizing data ingestion, training, and deployment — experiment setup time down 85% with 5x throughput across 10 engineers; mentored 2 junior engineers to feature ownership.

### ViTAA Medical Solutions | SOFTWARE DEVELOPMENT ENGINEER I

Montreal, Canada | June 2022 — June 2024

- Defined and implemented core data structures and coding standards across the codebase, cutting new-engineer ramp-up time from 3 weeks to 1 and reducing PR review time and review cycles by over 50%.
- Designed custom **PyTorch** UNets for 3D cardiovascular auto-segmentations across four domains, achieving a 92.84% Dice score and reducing QA manual adjustments by over 27%.
- Developed an automated CI/CD pipeline using **AWS CodeDeploy, EC2, and S3**, reducing test execution times by 43%; established a unit-testing framework (300+ tests) and trained 10+ engineers.

### Celestica | PRODUCT DATA ANALYST

Toronto, Canada | May 2019 — June 2021

- Built **VBA** automation and a neural-network component-cost model (92% accuracy), saving 20+ hours weekly and supporting \$1.5M in annual savings via **Tableau, Power BI, and Excel** analytics.

## Projects

---

### Coda | AI COMPLIANCE AUTOMATION

2026

- Built a suite of **Claude Code** skills that statically inspect repositories and auto-generate auditor-ready documentation across 3 compliance frameworks (**SOC 2, HIPAA, PCI-DSS**).
- Mapped application code (auth, data flows, cryptography) to control criteria, producing evidence-backed control narratives across 5 language stacks (Express, Django, Go, Spring, Rails).

### Satellite | MULTI-AGENT INVESTMENT ANALYSIS

2026

- Designed a self-contained **Claude Skill** orchestrating 11 specialized subagents across 2 strategy engines with progressive-disclosure architecture, screening 20–30 candidates to 8–12 finalists per run.
- Engineered a 3-tier degradation-tolerant data layer with per-field source tagging, plus a pass/fail validation harness (hard gates, weighted scoring) that raised output rubric scores from 69 to 86.

## Leadership & Community Involvement

---

### Brampton Robotics | ENGINEERING JUDGE

2022 — Present

- Evaluated regional and provincial robotics teams on engineering design and system performance; previously led a 15-engineer team to consecutive national VEXU championships (2018–2019), owning an autonomous **C++** control system.

## Education & Certifications

---

### PMI | PROJECT MANAGEMENT PROFESSIONAL (PMP)

2026

### DeepLearning.AI | DEEP LEARNING SPECIALIZATION

2022

### Toronto Metropolitan University | B.ENG MECHATRONICS ENGINEERING

2021

## Technical Skills

---

**Languages** Python, Java, JavaScript, TypeScript, C/C++, SQL, Bash

**Frameworks** PyTorch, TensorFlow, Scikit-Learn, Pandas, NumPy, React.js, Next.js, Node.js, Django

**Cloud & DevOps** AWS (S3, EC2, Lambda, SageMaker), GCP, Docker, Kubernetes, Slurm, Git, Unix/Linux, CI/CD

**AI / Agentic** Claude Code, MCP, multi-agent orchestration, RAG, prompt & context engineering, LLM evaluation