

HANGYU FENG

Seattle, US · (425) 736-5398 · fhy@uw.edu · GitHub · LinkedIn

SUMMARY OF QUALIFICATIONS

- 3+ years at Intel as a deep-learning software engineer building **FPGA inference infrastructure**: OpenVINO model benchmarking, **INT8 calibration**, C++ post-compilation tooling, reproducible CI/CD, Docker packaging, and static-analysis gates
- MS in Applied & Computational Mathematics** at UW (GPA: 3.90/4.0), with graduate work in high-performance scientific computing, numerical methods, AI robotics, randomized algorithms, and trading systems
- AI2 / UW research fine-tuning and evaluating **VLA foundation models** (MolmoAct2, $\pi_{0.5}$, SmolVLA) across real-robot and simulated manipulation settings

SKILLS

Languages & Tools: Python, C++, Bash, Ruby; PyTorch, OpenVINO, NumPy, SciPy; Git, Docker, Linux

PROFESSIONAL EXPERIENCE

Graduate Student Researcher

February 2026 – April 2026

Allen Institute for AI (AI2)

Seattle, WA

- Fine-tuning and evaluating **MolmoAct2**, a **Vision-Language-Action foundation model**, on manipulation tasks; designed task-success evaluations across simulation benchmarks and zero-shot real-world scenarios under Jiafei Duan and Prof. Dieter Fox
- Curated teleoperation datasets for physical-robot deployments; contributed to end-to-end training, simulation, and inference workflows

Deep Learning Software Engineer

January 2022 – June 2025

Intel Corporation

Toronto, Canada

- Benchmarked throughput and latency of 10+ CNN variants compiled via OpenVINO for FPGA inference, including **INT8 calibration** via the OpenVINO PTQ API; profiling reports informed two silicon architecture revision cycles
- Designed and implemented **AOT-Splitter in C++**, reverse-engineering OpenVINO AOT artifacts into memory payloads and configuration blobs for bare-metal FPGA provisioning without the full OpenVINO runtime
- Owned daily, weekly, and monthly **CI/CD pipelines**: reproducible builds with pinned dependencies / content-addressed artifacts and Python QOR reports (Excel/PDF/HTML) across build configurations
- Integrated **Coverity static analysis** as an advisory-to-blocking CI gate; shipped the official Docker image, collapsing install docs into a `docker pull`; built Debian-packaging tooling
- Led weekly paper-reading sessions on quantization, pruning, inference architectures, and LLM / diffusion research; created the team's on-device ML forum

PEY Software Developer

June 2019 – September 2020

Financeit

Toronto, Canada

- Launched the Centah rebuild in Ruby on Rails / Ember.js; designed and shipped the Financeit/Centah REST API integration between two fintech platforms
- Built unit, integration, and end-to-end test suites; implemented anti-fraud systems across both platforms

EDUCATION

Master of Science in Applied & Computational Mathematics | GPA: 3.90/4.0

University of Washington

Seattle, WA

- Relevant Coursework:** Probability & Random Processes, Applied Complex Analysis, Scientific Computing (4.0), ODE/PDE Methods, **High-Performance Scientific Computing**, Randomized Algorithms, **AI Robotics (CSE571, 4.0)**, Trading Systems (4.0)

Honours BSc, Computer Science Major (AI focus) + Mathematics Major | with Distinction

University of Toronto

Toronto, Canada

- Math:** Real Analysis II, Abstract Mathematics, Groups & Symmetry, Galois Theory, Mathematical Logic, Probability & Statistics, Chaos & Dynamical Systems; **CS:** AI, Machine Learning, Computer Vision, Computational Linguistics

PROJECT EXPERIENCE

VLA Fine-tuning vs. ACT Policy on Robot Arm | CSE 571 (Prof. Dieter Fox)

- Fine-tuned $\pi_{0.5}$ and **SmolVLA** (450M) on a curated 50-demo SO-101 dataset; benchmarked against ACT (100% success across 20 trials, mean radial error 3.46 ± 2.86 cm); trained on UW Hyak multi-GPU cluster; dataset on HuggingFace

Stochastic PDE Numerical Methods | AMATH 581 (Prof. Eli Shlizerman)

- Spectral/FEM PDE solvers: 3-D Gross-Pitaevskii equation on a 32^3 grid (FFT Laplacian, `scipy` ODE integration); FEM and finite-difference SPDE solvers under Euler-Maruyama and Milstein time-stepping

Social Distance Detection | Introduction to Image Understanding (UofT)

- Built a video inference pipeline combining RPN / Mask R-CNN with EMD loss, background subtraction, and perspective homography to map pedestrian detections into calibrated metric distances for violation flagging