Duong V. Nguyen

duongnviet1708@gmail.com
duongnviet.com

PROFILE

Edge AI engineer at VinRobotics, doing research with Dr. An Thai Le on machine learning on the hardware it actually runs on: how to keep a model accurate when its hardware is noisy or drifting, and how to evaluate it so the result holds outside the lab. First author of two papers under review and co-author of a journal paper on single-lead ECG.

EDUCATION

College of Engineering & Computer Science, VinUniversity — Ha Noi, Viet Nam
B.Sc. in Electrical and Computer Engineering — 2021 – 2025

- English-medium program validated by **Cornell University**; Dean's List, 2025.
- Capstone: InMoovXtra, a speech-interactive humanoid robot (see Projects).

Phan Boi Chau High School for the Gifted — Nghe An, Viet Nam
Physics specialized class — 2018 – 2021

PUBLICATIONS

1. **Duong V. Nguyen** and An Thai Le. *Budget-Aware Step Selection for Forward-Only Recovery of Analog Accelerators*. Under review.
2. **Duong V. Nguyen***, M. Vu* and An Thai Le. *Beyond Collision AUC: A Paired CARLA Benchmark for Event-Based Collision Warning*. Under review. (*equal contribution)
3. H. X. Nguyen, **Duong V. Nguyen**, H. H. Pham and C. D. Do. *MPCNN: A Novel Matrix Profile Approach for CNN-Based Single Lead Sleep Apnea in Classification Problem*. IEEE Journal of Biomedical and Health Informatics, 2024. *Paper*.

RESEARCH EXPERIENCE

VinRobotics | Edge AI Engineer | Ha Noi, Viet Nam — 2025.08 – now

Forward-only recovery of drifting analog accelerators — first author

- Built experiment software for simulated analog, photonic and spiking accelerators (AIHWKit, pytorch-onn, Rockpool); in a controlled testbed, single-step descent boundaries exceed whole-run failure boundaries by 27.2× (geometric mean), so early progress is a poor guide to a full run.
- Derived a finite-horizon bound on noise-driven drift and used its step scale to pick a learning rate from donor devices, without searching on the target device.
- On 36 held-out simulated crossbars, the selector needed 93.62% fewer forward evaluations than searching on the target, at a 0.399-point accuracy cost; the paper argues for reporting accuracy and configuration cost together.

A paired CARLA benchmark for event-based collision warning — co-first author

- Showed that on an existing event-camera corpus, episode duration alone reaches collision AUC 0.980, so the standard score can reward a shortcut instead of detection.
- Built an event-camera data pipeline in CARLA and 440 collision/near-miss pairs, covering stationary and moving vehicles and pedestrians, in which each pair shares scene and camera conditions, so a detector must separate contact from a matched near miss.
- Spiking and temporal-convolution detectors reach test AUCs of 0.60 and 0.57 but pair concordance of 69.8% and 78.6%, so the two measures rank them in opposite orders; thresholds set for 90% collision recall still raise frequent near-miss alarms.

Both studies use evaluation built to be audited: paired tests, grouped bootstrap, multiple-comparison correction, and provenance for every reported number.

VinUni-Illinois Smart Health Center | Undergraduate Researcher — 2023.08 – 2024.08

- Single-lead ECG sleep-apnea detection with Dr. Cuong Do and Dr. Hieu Pham: distance profiles over the whole heartbeat, instead of R-peaks alone, fed to lightweight CNNs (MPCNN, IEEE JBHI 2024, second author).
- 92.11% per-segment accuracy on PhysioNet Apnea-ECG and 81.25% on UCDDB; the same inputs improved existing small detectors (0.16–0.29 MB) with their architectures unchanged.

ENGINEERING EXPERIENCE

VinRobotics | Edge AI Engineer | Ha Noi, Viet Nam 2025.08 – now

Edge AI on NXP i.MX93 and Qualcomm RB8

- Brought up the NXP FRDM-i.MX93 on embedded Linux: device-tree changes, a MIPI CSI-2 camera driver ported into NXP's kernel, and a real-time (PREEMPT_RT) image built with Yocto.
- Packaged NCNN with Yocto recipes and a cross-compilation SDK, and deployed a PyTorch control policy that runs in about 0.3 ms per step on one Cortex-A55 core.
- Brought up the Qualcomm RB8 with Qualcomm Linux and Ubuntu, and benchmarked a humanoid robot's control policy on its CPU for inference time, CPU load and latency.

VinRobotics | Intern | Ha Noi, Viet Nam 2025.02 – 2025.08

- Real-time firmware on Teensy 4.1 and Zephyr RTOS for a 24-DOF humanoid robot head (lip sync, eye motion, facial expressions), commanded from ROS 2.
- Tracked down motion conflicts, latency and electrical noise with signal-level testing.

ARIS | Technical Lead, Unity/Vuforia AR app for medical education 2024.02 – 2024.08

SELECTED PROJECTS

Stewart Platform Neck for the Unitree G1 Humanoid 2025

- Designed the electrical system: a 6 V, 10 A servo supply and PCA9685 PWM from an Arduino Uno. Decoupling, separate power and signal wiring and split grounds cut PWM jitter from $\pm 8\,\mu$s to $\pm 1.5\,\mu$s and ended resets under current peaks above 7 A.
- Developed the face tracking and control: Haar detection with CSRT tracking, undistorted yaw and pitch, and inverse kinematics for six servos at 50 Hz; the head follows a face at 30 FPS with 48.6 ms end-to-end latency.

4-DOF Arm with Computer Vision 2024.09 – 2025.01

- Block detection on a Jetson Nano: HSV masks filtered by size, shape and work area; IoU 0.92–0.96 and F1 0.92–0.97 across three colours and two block sizes.
- Camera calibration from 44 chessboard views and a solvePnP pixel-to-world transform with 10.2 px mean reprojection error (about 0.5% of the image width); pick-and-place succeeded in 77–84% of trials.

InMoovXtra: Speech-Interactive Humanoid Robot 2024.08 – 2025.06

Capstone, VinUniversity

- Power and servo control for a life-sized humanoid built on the open-source InMoov platform: a 6 V servo supply, Arduino Mega and PCA9685 boards over I^2C, and every servo's channel, direction and limits set in MyRobotLab.
- Mapped voice commands to gestures, facial expressions and pick-and-place routines; the robot answered in about 2 s and changed expression in under 1 s.

Video demos of these projects are at duongnviet.com.

AWARDS AND HONORS

Dean's List, VinUniversity 2025

Finalist, Eureka Scientific Research Student Award 2024

Third Prize, provincial Physics merit competition 2020.12

SKILLS

Machine learning: Python, PyTorch, NCNN, OpenCV; AIHWKit, pytorch-onn, Rockpool; CARLA.

Embedded and robotics: C/C++, Rust, Zephyr RTOS, embedded Linux and Yocto, ROS 2; NXP i.MX93, Qualcomm RB8, Jetson Nano, Teensy and Arduino.