

Minh Phan

MINNEAPOLIS, MINNESOTA
VIETNAMESE + ENGLISH

Computer Science student building reliable software, ML systems, local AI applications, and computer vision experiments.

quangminhph07@gmail.com · github.com/Harryphan72007 · harryphan72007.github.io/minh-phan-portfolio

EDUCATION

B.S. Computer Science

EXPECTED MAY 2028
GPA 3.93 / 4.00

University of Minnesota Twin Cities

- Focus: computer vision, vision robustness, efficient deep learning, parameter-efficient fine-tuning, and AI systems.

RESEARCH AND EXPERIENCE

Software Engineering Intern

MAY 2026 - PRESENT
HO CHI MINH CITY, VN

FPT Software

- Built Mega-ASR Studio with FastAPI, React, Vite, and reusable local ASR inference.
- Implemented batch and single-file APIs, inference locking, validation, recording, previews, system status, and decode/memory error handling.

Undergraduate Research Volunteer

MAY 2026 - PRESENT
MINNEAPOLIS, MN

Ding Lab - University of Minnesota Twin Cities

- Prepared datasets, reviewed preprocessing outputs, and flagged inconsistent or low-quality samples for a PhD-led AI hardware and systems project.
- Ran and analyzed computer-vision experiments while supporting reproducible research workflows.

RESEARCH WORK

The Shape of Noise

WORKSHOP CO-AUTHOR
COMPUTER VISION

Layer-Wise Perturbation Profiles for Diagnosing Vision Robustness - accepted as a poster to the CTB workshop at ICML 2026

- Co-authored with S. Nguyen, V. G. Bao, and T. P. Le; ran the controlled ResNet-50 and ConvNeXt-Tiny experiments on CIFAR-10 and CIFAR-10-C.
- Compared full fine-tuning, LoRA, and profile-selected layer subsets across parameter counts and multiple seeds.
- Reported in the paper: Top-k LoRA on ConvNeXt-Tiny reached 92.57% corrupted accuracy with 0.04M trainable parameters, against 94.95% with 27.83M for full fine-tuning.

SELECTED PROJECTS

Aerial Object Detection Benchmark

ML SYSTEMS - INFRASTRUCTURE BUILT, GPU RUNS
PENDING

Single-protocol VisDrone comparison across Faster R-CNN (ResNet-50, Swin-T, VMamba-T) and RT-DETRv2-L. Built the COCO conversion, resumable Optuna studies, checkpoint lifecycle, GPU adapter smoke gate, and schema-validated result bundles. 498 tests across 67 modules; no benchmark result is claimed yet.

[Python](#) / [PyTorch](#) / [MMDetection](#) / [Optuna](#) / [VisDrone2019](#) / [COCO](#)
github.com/Harryphan72007/aerial-object-detection-benchmark

NoteFlow AI

LOCAL-FIRST AI - RELEASED V0.1.0

Turns notes, audio, and scans into reviewable records that keep source, model output, and human correction separate. Built the FastAPI system of record, correctable ASR/OCR records, WER/CER and numeric-mismatch comparison, audit history, and the React review interface. 19 backend tests plus frontend gates in CI.

[FastAPI](#) / [SQLAlchemy](#) / [React](#) / [TypeScript](#) / [ASR](#) / [OCR](#) / [Ollama](#)
github.com/Harryphan72007/NoteFlow-AI

TECHNICAL SKILLS

LANGUAGES Python, TypeScript, Java

ML + VISION PyTorch, MMDetection, Optuna, LoRA, fine-tuning, COCO evaluation, multi-seed analysis

AI APPLICATIONS Local ASR, OCR, Ollama, human-review and document-processing workflows

BACKEND + WEB FastAPI, SQLAlchemy, Alembic, REST APIs, React, Vite

TOOLING Git, GitHub Actions, pytest, Docker, LaTeX