

Đoàn Việt Hưng

☎ (+84) 394 529 869 📍 Tay Thanh Ward, HCMC ✉ hungdoan29102003@gmail.com
🌐 linkedin.com/in/hung-đoàn-a636502ba 📄 github.com/DoanVietHung03 🌐 hungdv.io.vn

PROFILE SUMMARY

Computer vision engineer who takes models from experiment to on-premise deployment: face recognition, face anti-spoofing, license plate recognition, PPE detection and camera-based analytics on multi-stream RTSP and NVIDIA GPU edge devices. Optimizes inference with ONNX Runtime and CUDA and ships with Docker.

WORK EXPERIENCE

HPT Vietnam Corporation

company link [🔗](#)

Intern || HSV Department

July 2025 - Nov 2025

AI - CV Solution Engineer (Collaborator) || HSV Department

Dec 2025 - May 2026

AI - CV Solution Engineer (Fresher) || HSV Department

June 2026 - Present

- Contributed to the company's vision AI platform for security and operations (face recognition, face anti-spoofing, license plate recognition, PPE detection); co-author on the platform's intellectual property certificate.
- Designed and deployed AI surveillance POC systems for enterprise clients on multi-stream RTSP (H.264) feeds, running on NVIDIA GPU edge devices.
- Optimized real-time inference pipelines with ONNX Runtime and CUDA, keeping multi-stream processing stable under hardware constraints.
- Containerized models and backend services with Docker for reproducible on-premise deployment at enterprise clients, including major food and beverage brands.
- Translated business requirements into system-level AI architectures, balancing accuracy and latency for each use case.

PROJECTS

Surface Defect Inspection with a Line-Scan Camera (Anomaly Detection)

portfolio (private source) [🔗](#)

July 2026 - Oct 2026

- Few-shot anomaly detector for a line-scan inspection line: PatchCore on MobileNetV3-Large (FP16) learns from only 10 good images and needs no defect images; each 1024×1792 frame is tiled, scored and thresholded on 50 separate good images, returning OK/NOK, an anomaly map and boxes.
- 92.3% recall (24/26) with 2.8% false alarms (3/107) on a 133-image test set; 43 ms P50 latency on an RTX A4000 (about 683 ms on CPU with OpenVINO), within the 50 ms per-image budget.
- Operations dashboard (aiohttp, React 19 + TypeScript, SQLite, Basler camera via pypylon) where operators train and activate a new product from the browser; few-shot evaluation across 8 seeds with nested leave-one-out validation.

Internal Smart Attendance System (Face Recognition)

private source [🔗](#)

August 2025 - Dec 2025

- **Phase 1 (self-developed):** two-stage face recognition with an iResNet embedding gallery and Faiss search, YOLOv11n-face / RetinaFace detection and ByteTrack tracking; benchmarked detectors on accuracy-latency trade-offs under GTX 1650Ti constraints.
- Built the FastAPI + MySQL backend and real-time web UI, containerized with Docker Compose; accelerated inference with ONNX Runtime and CUDA.
- **Phase 2 (Frigate v0.16 NVR):** customized Frigate with RetinaFace and InsightFace (ArcFace) on ONNX

Runtime / NVIDIA GPU for multiple RTSP streams; extended its event pipeline for face-based access logging via MQTT and a web dashboard.

Real-time POS and Cashier Monitoring System

[github link](#) 

Dec 2025 - Jan 2026

- Real-time cashier fraud detection combining deep learning and classical CV: MediaPipe hand tracking and a custom TensorFlow Lite drawer-state classifier.
- Finite State Machine for valid transaction flows, flagging anomalies such as "Void after Cash Access" and "Open without Transaction"; configurable time-window matching of camera events against POS logs.
- Motion gating (background subtraction) to cut idle CPU, anti-tampering alarms for a covered or obstructed camera, and a multi-threaded capture / inference / UI (Streamlit) pipeline for high FPS and low latency.

Vision Measurement Toolkit: Distance and Speed from a Single Camera

[github link](#) 

Jan 2026 - Present

- Distance estimation from a monocular IP camera with two approaches: geometric (IPM + homography, Lucas-Kanade optical flow to stabilize calibration, YOLO) and deep learning (Depth-Anything-V2 with a 3D pinhole model), plus intrinsic / distortion calibration.
- Multi-threaded design decoupling AI inference from UI: 49 FPS at 13.34% MAPE (geometric) and >230 FPS display at 15.07% MAPE (deep learning); analyzed trade-offs between latency, robustness, hardware and accuracy.
- Added lane-based vehicle speed with ground-truth validation, and merged two earlier experiment repositories into one package with unit tests.

SKILLS

- **Programming Languages:** Python, C++.
- **Computer Vision / ML:** PyTorch, TensorFlow / TFLite, Ultralytics (YOLO), OpenCV, MediaPipe, InsightFace, ByteTrack, Depth Anything, anomaly detection (PatchCore).
- **Deployment:** ONNX Runtime, CUDA, OpenVINO, Docker / Docker Compose, RTSP (H.264), NVIDIA GPU edge devices.
- **Backend / Data:** FastAPI, MySQL, SQL Server, MQTT, Faiss, NumPy, Pandas.
- **Tools:** Linux, Git, Jupyter Notebook, Streamlit.
- **Languages:** English (Professional working proficiency).

EDUCATION

VNU-HCM UNIVERSITY OF SCIENCE

Oct 2021 – Oct 2025

BS in Computer Science

- Major in Computer Vision
- GPA: 3.0/4.0