

Benedict Chan

+65 8468 7191 | benedictchan247@gmail.com | linkedin.com/in/benedictchannn | github.com/BenedictChannn
Singapore citizen — **H-1B1-eligible (U.S.)**; open to relocation

EDUCATION

National University of Singapore

Aug 2025 – Present

Master of Science in Robotics

National University of Singapore

Aug 2022 – May 2025

Bachelor of Engineering in Computer Engineering, Minor in Innovation & Design Program

- Accolades: NUS Engineering Scholars Program (accelerated B.Eng + M.Sc), DSTA TechUP (Pioneer cohort)
- Relevant Coursework: Image Processing and Analysis, Robot Perception, Data Structures and Algorithms, Signals and Systems, Software Engineering and OOP, Real-Time Operating Systems, Discrete Structures

TECHNICAL SKILLS

Programming: Python, TypeScript, C/C++

ML & Vision: PyTorch, OpenCV; CLIP/ViT; YOLO; LVLMs; scikit-learn, XGBoost; Hydra

Agents: LangGraph; tool/JSON-schema calling; routing & consensus; embeddings & vector search (Weaviate/HNSW)

Simulation & Robotics: ROS, Blender/SMPL-X, UE4, CARLA; LiDAR point-clouds; SLAM/ESVO (C++); FFmpeg

Systems & Cloud: Linux, Docker, Git/GitHub Actions; REST APIs; GCP Tools – Firebase/Firestore, GCS, Vertex AI

WORK EXPERIENCES

Founding Software Engineer

May 2025 – Dec 2025

GrowtricsAI (AI EdTech Startup)

Singapore

- Architected a **document-parsing pipeline** to replace manual labeling, converting unstructured exam papers into structured Q&A data, which achieved **83% conversion rate**, accurately populating **10k+ questions**.
- Developed a **video generation pipeline** using LangGraph + Flutter that transforms students' graded performance into personalised, step-by-step **explainer videos** to address student's misconceptions.
- Defined **SLOs** and cost projections for GenAI workloads; owned entire document-parsing features roadmap to raise extraction quality and **extend coverage** across subjects, curricula, languages, and modalities.

Generative AI Researcher Intern

Dec 2024 – Sep 2025

SIT × NVIDIA AI Center

Singapore

- [MANUSCRIPT IN PREPARATION] *SynDORBench: Benchmarking Human Detection and Action Recognition with DORI-Calibrated Distances*.
- Engineered a **scalable synthetic-data engine** using open-source datasets (SMPL-X, AMASS, and BABEL) and implemented an auto-labeling mechanism (pose/bbox/mask), which accelerates CV/VLM data collection.
- Developed a **DORI-aligned surveillance dataset** (200k+ images & 630k+ QA pairs) with controlled **context parameters** to enable **robust benchmarking** of presence & action recognition models.
- Built a **unified evaluation framework** for human-presence & action recognition, with Hydra configs and Docker/CUDA to benchmark **YOLO** models against compact **LVLMs**.

Algorithm Engineer Intern

Feb 2024 – Jul 2024

Outsight (3D Spatial Intelligence Startup)

Paris

- Improved perception validation accuracy by **67%** in large-scale environments through **high-fidelity People Flow Monitoring** simulations with realistic crowd dynamics and stochastic destination assignments (UE4/CARLA).
- Supported pre-deployment validation of over **300 LiDARs** across multiple terminals at a major SEA airport.
- Cut simulation setup time by over **90%** via an **end-to-end scenario tooling suite** for scenario configuration, point-cloud dataset generation, and fixed-seed regression checks, accelerating iteration velocity for algorithm team.

Computer Vision Engineer Intern

May 2023 – Aug 2023

DSO National Laboratories

Singapore

- Increased drone navigation accuracy by **20%** in low-light, high-speed conditions by developing an **Event-based Stereo Visual Odometry (ESVO)** system using **C++**
- Achieved **5x latency reduction** in state estimation (80%↓) by re-implementing non-linear solvers with **SymForce symbolic computations**, enabling real-time control feedback in drone navigation systems.