Arjun Prakash

+65 8286 7206 | arjun024@e.ntu.edu.sg | arjunprax.com
github.com/ArjunPrax | linkedin.com/in/arjun-prakash2006

Education

Nanyang Technological University 2024 – 2028
Bachelor of Engineering in Electrical and Electronics Engineering Singapore
Second Major in Data Analytics GPA: 4.01/5.00

Experience

Panasonic R&D Singapore Jun 2026 – Present
Machine Learning & Perception Engineering Intern *Singapore*

- Built RGB-D perception, camera-calibration and trajectory-auditing tools for learned-policy evaluation.
- Integrated VLA and imitation-learning policies (π0.5, GR00T N1, ACT, Diffusion Policy) with sensing and teleoperation via ROS 2 and MoveIt 2 Servo across UR5/UR5e, Unitree G1 and RoboParty V1.
- Trained and evaluated imitation-learning and PPO policies in Isaac Lab; developed motion retargeting and diagnosed execution using trajectories, logs and physical testing.
- Built Docker-based Isaac Sim/Isaac Lab workflows with a shared simulation/hardware control path; cross-validated Isaac PhysX and MuJoCo models against real torque measurements.
- Developing geometry-conditioned ACT for in-hand reorientation using CAD-free RGB-D geometry and gripper state; manuscript in preparation for IROS 2027 with NTU MARS Lab.

Nanyang Technological University Jul 2026 – Present
Undergraduate Researcher – ML Systems & 3D Perception *Singapore*

- Implemented voxelized LiDAR initialization for Gaussian semantic occupancy, with optional density-based opacity priors and an unchanged 28-D interface to the downstream encoder and decoder.
- Built a mask-aware occupancy evaluator with strict checkpoint loading, per-class IoU, GPU latency and peak-memory reporting to support controlled camera–LiDAR comparisons.
- Designing confidence-gated temporal Gaussian updates to reuse reliable geometry and prioritize uncertain or changed regions, aiming to reduce latency while preserving occupancy accuracy; targeting ICCV 2027.

Nanyang Technological University Jul 2025 – Jun 2026
Undergraduate Researcher – ML Inference & Computer Vision *Singapore*

- Applied INT8 post-training quantization with under 0.4 percentage-point accuracy loss; recalibrated confidence thresholds to recover 93% defect recall without retraining.
- Built a custom Verilog convolution accelerator on Xilinx ZCU104, achieving 3.3× faster inference and 6.3× lower energy per image than the onboard ARM CPU.
- Trained a ResNet18 PCB-defect classifier to 85.3% accuracy and adapted an SSD300 ADAS detector to 44.6% mAP@0.5 across four safety-critical classes; presented at IEEE ICIEA 2026.

Golden Horizon Sep 2025 – Dec 2025
Product Owner – AI Analytics Platform

- Led development of an AI analytics platform combining market data, event signals and LLM-generated insights for trading and portfolio decision support.
- Built event-to-price data pipelines and translated analysis of 20+ competitors into product positioning and execution priorities.

Projects

Compact Vision Inference Service

- Deployed an ONNX vision model through a containerized API; INT8 yielded an approximately 4× smaller model and hardware-dependent 1.5–3× CPU speedups with under one percentage point of accuracy loss.

Semantic Search over Robot Videos

- Built timestamped semantic search over robot videos; achieved approximately 60–80% Recall@5 on visible-state queries with millisecond-scale retrieval.

Task-Aware Joint-Space Safety Envelope Implemented

- Built a task-specific runtime guardrail from recorded joint angles and speeds that rejects out-of-envelope commands, aborts execution and returns explicit diagnostics.

Leadership

MLDA@EEE – Main Committee, Business Development Aug 2025 – Present
Managed NVIDIA, Micron and industry partnerships; led teams delivering applied-ML projects and technical showcases.

Technical Skills

ML and Optimization: PyTorch, TensorFlow, scikit-learn, transfer learning, ONNX export, INT8 PTQ
Inference and Deployment: Containerized inference, REST APIs, latency/memory benchmarking, Docker, Linux
Multimodal Retrieval: Visual feature extraction, semantic search, vector indexing, LangChain, RAG
Perception and Robotics: RGB-D/LiDAR, Gaussian occupancy, ROS 2, MoveIt 2 Servo, Isaac Sim/Lab, MuJoCo
Programming and Data: Python, C/C++, SQL, NumPy, pandas, Git

Publication

A. Prakash and X. S. Loo, "FPGA-Accelerated Edge Vision: PCB Defect Detection and ADAS Object Detection on Xilinx ZCU104," IEEE ICIEA 2026, Catania, Italy. Presented.