

Yugendren S/O Sooriya Moorthi

+65 9106 2034 | yugen@mericanii.com | github.com/yugendren | linkedin.com/in/yugendren | Singapore

Education

Singapore Institute of Technology <i>BEng (Hons) Robotics Systems cGPA 4.75/5.0 Provost's Honour List (every eligible semester)</i>	Sep 2022 – Apr 2026 Singapore
Temasek Polytechnic <i>Diploma in Aerospace Engineering cGPA 3.29/4.0</i>	Apr 2017 – Mar 2020 Singapore

Experience

Government Technology Agency (GovTech) <i>Software Engineer Intern — Production ML Infrastructure</i>	Dec 2024 – Present Singapore
---	---------------------------------

- Led ML ablation studies and retraining cycles for RoBERTa and Random Forest ensembles, analyzing production data to optimize precision and recall for scam detection.
- Developed a Gemini-powered GenAI wrapper to catch harder-to-detect scam sites (e-commerce, crypto) that classical ML missed, improving F1 by up to 30% while reducing false positives.
- Extended the Recursive Machine Site Evaluator (RMSE) to resolve domain ownership via per-TLD WHOIS / RDAP services, improving extraction coverage by about 20%.
- Built and deployed an AWS Lambda blacklisting service in the production stack for real-time domain enforcement, enabling faster decisioning in scam-mitigation workflows.
- Engineered schema-aware ETL pipelines consolidating high-volume data from multiple backend services (MongoDB, DynamoDB) into S3 and Aurora RDS to power GASP (Government Anti-Scam Products); validated throughput via Athena load testing.

SIT-NVIDIA Joint Lab <i>AI Research Intern — Multimodal Video Embeddings</i>	Jun 2024 – Dec 2024 Singapore
--	----------------------------------

- Ran experiments on attention mechanisms (self, multi-head, spatial, temporal, cross) for composite video embeddings in Multimodal Large Language Models.
- Implemented composite embedding strategies in the vision tower, supporting PhD-led research on vision-tower ablations.
- Set up infrastructure to train, fine-tune, and validate Multimodal LLMs on a single NVIDIA A6000 GPU.
- Implemented key-frame selection to sample the most informative frames from video instead of processing every frame, reducing redundant compute in the embedding pipeline.

Projects

Orphan — Multi-Agent RL for Pursuit-and-Intercept <i>Isaac Sim, JAX, MAPPO, self-play, ONNX</i>	Dec 2025
--	----------

- Trained a multi-agent RL pursuer via MAPPO self-play between co-adapting agents for autonomous pursuit-and-intercept in GPS-degraded, link-limited environments; converged from 0% to 76–88% intercept success within 100 iterations against a co-trained adversary, with no hand-coded motion-planning heuristics.
- Prototyped in NVIDIA Isaac Sim (USD, UsdPhysics), then built a custom JAX physics env (jit / vmap / lax.scan) with velocity-command control and domain randomization; trained on cloud GPUs (RunPod, AWS EC2 / SageMaker) and exported via ONNX Runtime for sub-1ms edge inference. Simulation-only.

Prometheus — Cognitive-Radio RL <i>JAX, MAPPO, population-based self-play, Elo, MARL, ONNX</i>	Feb 2025
---	----------

- Developed an RL agent learning adaptive frequency-hopping policies under heavy interference and spectrum congestion; built a custom JAX RF physics engine (256-bin spectrum, multipath fading, SINR / AGC / BER) running at tens of thousands of env-steps/sec.
- Trained via MAPPO + population-based self-play (3×3 Elo-sampled populations, 25% scripted-adversary injection) on cloud GPUs (RunPod, AWS EC2 / SageMaker) across five interference archetypes; a 2×512 MLP (968K params) achieved 100/100/80/40/0% packet-delivery across progressively harder scenarios. Simulation-only.

Skills

Machine Learning & RL: PyTorch, JAX (jit, vmap, lax.scan), TensorFlow, NumPy, ONNX, OpenCV, attention mechanisms, Multimodal LLMs, GenAI (Gemini), MAPPO, MARL, self-play, population-based training, Elo, A2C / PPO / SAC, MuJoCo, domain randomization, sim-to-real, RoBERTa, Random Forest

Programming Languages: Python, C++, C, SQL, Bash

Simulation & Systems: NVIDIA Isaac Sim, Omniverse, USD / UsdPhysics, ROS 1/2, Gazebo, RViz, Linux

Cloud & Data: AWS (Lambda, EC2, Glue, Athena, DynamoDB, SageMaker, S3, Aurora RDS), RunPod, MongoDB, ETL, REST APIs, Docker, Git, CI/CD