Chua Han Chong

CHUA1163@e.ntu.edu.sg Singapore chuahanchong.github.io ChuaHanChong chuahanchong
ChuaHanChong

Professional Summary

- PhD student at MMLab@NTU working on embodied AI with spatial intelligence, through real-to-sim-to-real transfer and continual learning. Nine years of prior experience in data science, artificial intelligence, and MLOps, turning research methods into production systems at scale and leading cross-functional teams to deliver them.

Experience

Porolis Technologies Pte. Ltd., AI Engineer Lead
Singapore
May 2023 – Aug 2026
3 years 4 months

- Architected and deployed the production AI platform serving all gemstone analysis models (FastAPI, Redis, RQ, Docker, nginx, AWS EC2 and S3), covering API gateway and load balancing, distributed job queues with priority-based worker scheduling, GPU autoscaling on scheduled cost windows, fault tolerance (retries, dead-letter queues, fallback responses, exactly-once webhook delivery), and centralised logging and monitoring.
- Applied Tip-Adapter (training-free CLIP adaptation) and DualCoOp (dual prompt tuning for multi-label recognition) on the SigLIP2 vision-language encoder with only 560 annotations, reaching >75% binary accuracy for heat-treatment classification and >70% average overall F1 on 28-type inclusion recognition.
- Applied Segment Anything 2 with visual prompting for zero-shot gemstone masking and tracking in video, reaching J&F >85%.
- Deployed MiniMax-Remover to remove handlers and inpaint occluded regions in video, reaching >8.1 visual quality and >95% success rate as judged by Gemini 2.5 Pro under an LLM-as-a-judge protocol.
- Used Gemini 2.5 Pro with prompt chaining and reflection plus few-shot prompting for gemstone-versus-background detection (>85% average precision at IoU 0.75), and Gemini 2.5 Flash image generation for stain removal (DINOv2 similarity >90%).
- Raised colour-grading performance to 90% F1 via semi-supervised learning with high-confidence pseudo-labelling on unlabelled data, and applied Cleanlab confident learning to reduce label noise.
- Led adoption of Feast as a feature store and introduced MLflow Projects with a factory pattern, centralising feature management and standardising training execution for reproducibility and reuse.

Ideagen ICT Sdn. Bhd., Senior Machine Learning Engineer
Kuala Lumpur, Malaysia
Sep 2022 – Apr 2023
8 months

- Implemented and fine-tuned Microsoft's Table Transformer on in-house datasets, improving table parsing accuracy by at least 5%.
- Optimised the PDF parsing pipeline in Ideagen AuditXprt using AWS Textract with document layout analysis (pdfminer, deep layout parser), improving parsing accuracy by 40% and speed by 250%.
- Refined the OCR solution in AuditXprt using the ABBYY FineReader 12 engine.
- Used MLflow for experiment tracking and as a centralised model registry, improving deployment efficiency by 30% across ML and software engineering.
- Introduced Docker Compose for model inference APIs and Selenium regression testing for the PDF parsing tool, cutting debugging time by 30%.

Holmusk, Senior Data Scientist — Singapore
Jan 2022 – Aug 2022
8 months

- Directed a 20-person multi-disciplinary team to set the technical vision for an MVP, hitting each milestone two weeks ahead of schedule.
- Built hospitalisation and self-harm risk prediction models with a data enrichment method, reaching AUROC of at least 0.75.
- Engineered a calibration method for risk value adjustment across mental healthcare centres.
- Devised a decision-tree study protocol to differentiate mental health patient profiles for cohorts with limited features.
- Deployed MLflow as the team's MLOps platform and standardised post-training model evaluation reporting.

Holmusk, Data Scientist — Singapore
Mar 2019 – Dec 2021
2 years 10 months

- Launched customised Food Tracker AI into AIA Vitality across five countries, serving over 80k users and scoring over 2 million images.
- Developed lightweight transfer-learning models reaching >92% food detection and >90% food scoring accuracy.
- Implemented continual learning with elastic weight consolidation and rehearsal, enabling model refinement within 30 minutes post-launch.
- Pruned model weights with the TensorFlow Model Optimization Toolkit, cutting weight size by 10% with minimal accuracy loss.
- Pioneered an AWS SageMaker MLOps framework and doubled training throughput using multi-GPU EC2 instances.

Celcom Axiata Berhad, Data Scientist — Kuala Lumpur, Malaysia
Sep 2017 – Feb 2019
1 year 6 months

- Built customer lifecycle segmentation for monetisation and defection probability.
- Developed a cross-selling model over sequential purchase behaviour, exceeding campaign revenue growth targets by 5%.
- Built an uplift model for promotion response likelihood, cutting message cost by 10%.

Education

Doctor of Philosophy in Computing and Data Science — Singapore
Nanyang Technological University
Aug 2026 – present

- MMLab@NTU, affiliated with S-Lab. Advised by Professor Chen Change Loy.
- Supported by the NTU Research Scholarship.

Master of Science in Artificial Intelligence — Singapore
Nanyang Technological University
Aug 2023 – Jun 2026

- CGPA: 4.30/5.00.
- Coursework: AI6101 Introduction to AI & AI Ethics; AI6102 Machine Learning Methodologies & Applications; AI6103 Deep Learning & Applications; AI6104 Mathematics for Artificial Intelligence; AI6128 Urban Computing; AI6129 AI Master Project; AI6130 Large Language Models; AI6131 3D Deep Learning; AI6132 Generative AI for Visual Synthesis.

Bachelor of Engineering in Mechanical Engineering — Kuala Lumpur, Malaysia
Universiti Putra Malaysia
Sep 2013 – Jul 2017

- CGPA: 3.728/4.00.

Skills

Machine Learning and Research
Self-supervised and representation learning, vision-language models, continual learning, reinforcement learning, domain adaptation, transfer learning, model distillation, post-training calibration

Computer Vision
Image classification, object detection, semantic segmentation, visual prompting, zero-shot masking and tracking

NLP and LLMs
Prompt engineering, in-context learning, retrieval-augmented generation, semantic matching, named entity recognition, topic modelling

Statistics and Analysis
Hypothesis testing, A/B testing, multi-armed bandit testing, survival analysis

Languages and Libraries
Python, SQL, PyTorch, TensorFlow, scikit-learn, Hugging Face, OpenCV, pandas, NumPy, SciPy, Matplotlib, Seaborn

MLOps and Infrastructure
MLflow, Feast, TensorFlow Extended, FastAPI, Redis, Docker, AWS (SageMaker, EC2, S3), GCP (Vertex AI, Compute Engine), Optuna, Ax, Git

Languages

Chinese
Native speaker

English
Professional working proficiency

Interests

Research
Embodied AI with spatial intelligence, real-to-sim-to-real transfer, continual learning

Certificates

AWS Fundamentals Specialization Jul 2023
Amazon Web Services

Machine Learning Engineering for Production (MLOps) Specialization Jul 2023
DeepLearning.AI

Practical Data Science on the AWS Cloud Specialization Jul 2023
DeepLearning.AI and Amazon Web Services

Generative AI with Large Language Models Jul 2023
DeepLearning.AI and Amazon Web Services

Natural Language Processing Specialization Jun 2023
DeepLearning.AI

Mathematics for Machine Learning and Data Science Specialization 2023
DeepLearning.AI

AI for Medicine Specialization Jun 2020
DeepLearning.AI

Deep Learning Specialization Jun 2018
DeepLearning.AI

Machine Learning Specialization 2018
DeepLearning.AI

Projects

Learning Self-Supervised Visual Representations for Fine-grained Maritime Vessel Recognition

MSc dissertation, Nanyang Technological University.

- Curated over 3.1M vessel images across 21 fine-grained categories via a data-sourcing pipeline, deduplicated with a Self-Supervised Copy Detection model.
- Built supervised Vision Transformer baselines (ViT-B/L/H) and compared self-supervised paradigms including MAE, DINOv2, and I-JEPA.
- Addressed class imbalance with balanced sampling and logit-adjusted loss.
- Reached 98.82% test accuracy over 17 in-distribution vessel classes and 100% over 4 out-of-distribution classes, showing self-supervised pre-training generalises better than supervised-only baselines.

Application of the Bees Algorithm to Improve the Performance of Flywheel Energy Storage

BEng final year project, Universiti Putra Malaysia.

- Optimised composite flywheel rotor geometry with the Bees Algorithm under the Tsai-Hill failure criterion to maximise energy storage while minimising mass.
- Validated top designs with finite element analysis in ANSYS, improving usable energy by 23.5% and reducing rotor mass by 41.5% against the NASA G3 baseline.