Michael Hoon Yong Hau

+65 97769494 | michaelhoon@gmail.com | linkedin.com/in/michaelhoon | github.com/michael-hoon

Education

Singapore University of Technology and Design Sep. 2022 – May 2026
Bachelor of Engineering in Engineering Systems and Design, Minor in Artificial Intelligence (AI) *GPA: 4.59/5.00*

Chalmers University of Technology Jan. 2025 – Jun. 2025
Spring Exchange, Machine Learning, AI, Big Data Systems *Gothenburg, Sweden*

University of California, Berkeley Jun. 2023 – Aug. 2023
Summer Exchange, Econometrics *Berkeley, CA*

Experience

AI Engineer Intern Jun. 2025 – Sep. 2025
Government Technology Agency (GovTech) Singapore *Singapore*

- Identified inefficiencies in Redshift data warehouse (column encoding, workload management), migrating provisioned clusters to serverless architecture, reducing storage costs by ∼30% and query latency by ∼15%.
- Spearheaded an AWS SageMaker Unified Studio landing zone (Bedrock, Lake Formation, Glue, Athena, DataZone) with Terraform IaC, centralizing data access for 50+ users and cutting analytics onboarding time by ∼80%.
- Productionised an end-to-end data pipeline for clickstream analytics, ingesting 100M+ daily user events via Data Firehose, Glue, S3, Redshift Spectrum. Enabled analytics and dashboarding via SageMaker AI and Quicksight.

AI Engineer Intern Mar. 2024 – Dec. 2024
DSO National Laboratories *Singapore*

- Developed a full-stack local RAG pipeline for assisting researchers with internal workflows, reducing initial research and policy analysis by 10+ hours/week. Designed pipeline for Knowledge Graph generation from PDF files.
- Designed a containerized, Multi-Agent Graph-based RAG system that answered complex multi-hop queries, improving accuracy by ∼30% vs. standard RAG pipelines.
- Architected end-to-end RAG pipeline (LangGraph, Ollama), with Knowledge Graph database (Neo4j, MinIO S3, PostgreSQL), and LLM-as-a-Judge evaluation, with LangSmith for tracing and observability.

Data Analyst Intern Jan. 2024 – Apr. 2024
Poh Tiong Choon Logistics Ltd. *Singapore*

- Supply chain optimization, analyzed delivery driver datasets using Python/PowerBI to derive insights and improve scheduling of drivers, boosting driver utilization by ∼20%.
- Built scheduled Bash/Python ETL scripts to automate data cleaning, visualization, and analysis with PowerBI, slashing weekly report generation time from 8 hours to 1 hour.

Projects

LepakLah! | *Python, FastAPI, MySQL, Docker, Kubernetes* Aug. 2024 – Sep. 2024

- Delivered a Flutter + FastAPI cloud-native mobile app for a Dell public-good hackathon, reducing senior workshop activity planning time by ∼60%. 4.8/5 satisfaction, won 2nd Runner up representing University.
- Led backend development, deployed LLM system and Stable Diffusion workflow via Nvidia NIM to auto-generate personalized workshops. Containerised microservices with Docker/Kubernetes on Red Hat OpenShift, able to scale to support 500+ concurrent users.

Cloud-Based MRP System | *Python, Azure, Airflow, Docker* May. 2024 – Aug. 2024

- Designed backend system for Materials Requirement Planning (MRP) for a supply chain project, optimizing inventory management by forecasting material requirements and automating workflows.
- Built an end-to-end ETL pipeline with Python, SQL, Azure Cloud Platform, and orchestrated automated DAG workflows with Apache Airflow. Containerised back-end microservices with Docker.
- Integrated system with Azure SQL database to store and retrieve data on Microsoft Azure Cloud Platform, enabling scalable and secure cloud-based data storage.

Kaggle Transportation Analytics | *R, Machine Learning, XGBoost* Aug. 2024 – Aug. 2024

- Engineered a stacked ensemble (Mixed Logit + XGBoost) in *R*, predicting vehicle feature preferences with ∼88% accuracy, outperforming 30+ teams and winning 1st Runner up on internal University hackathon.
- Fine-tuned model hyperparameters and validated results on private test set with cross-entropy loss. Leveraged SHAP values for interpretability and identified key safety features driving user purchase decisions.

Skills

Languages: Python, SQL, R, Go
Tools: Git/GitLab, AWS Cloud, Linux/Bash, Docker, Kubernetes, Ollama, Terraform, Airflow, PostgreSQL, Neo4j
Frameworks: Pandas, NumPy, PySpark, LangGraph, HuggingFace, scikit-learn, FastAPI, tidyverse (R), Scrum

Awards

SUTD Honours and Research Programme Sep. 2022 - May. 2026
Scholarship *Singapore*

- Selected as one of three students out of a cohort of 500+ for the highly competitive SUTD Honours and Research Program (SHARP), securing a $15K research grant.

SUTD Global Merit Award Sep. 2022 - May. 2026
Scholarship *Singapore*

- Awarded full sponsorship (worth over $16K) under SUTD Global Leadership Programme for academic summer exchange at University of California, Berkeley.

A C Toh - SUTD Scholarship Sep. 2022 - May. 2026
Scholarship *Singapore*

- Awarded full sponsorship covering tuition fees, stipend, and on-campus housing (worth over $60K) for outstanding academic achievement.

KKH - SUTD Exchange Award Jan. 2025 - Jun. 2025
Award *Singapore*

- Awarded sponsorship worth $5K to support overseas semester exchange in Sweden.

IMDA SG Digital Young Leaders Programme Sep. 2024 - May. 2025
Talent Programme *Singapore*

- Selected as part of IMDA's talent development initiative for mentorship and networking through industry attachments.

Leadership

Secretary Mar. 2023 – Feb. 2024
SUTD Organisation of Autonomous Robotics (SOAR) *Singapore*

- Organised a university-wide robotics competition (RoboClash 2024), attracting over 100+ participants across 20+ teams, and judging coordination to ensure a successful large-scale event.
- Spearheaded sponsorship and partnership initiatives, successfully securing $50K+ in funding from external companies and university grants to support competitions, equipment upgrades, and outreach programs.
- Managed club operations, communications, and documentation, ensuring smooth coordination across project teams and executive committee members.