

WORK AUTHORISATION

Singapore Citizen, eligible for **H-1B1** (no lottery/100k fee, unlike H-1B), **J-1** US and **HPI** UK visas; **no petition required**.

EDUCATION

National University of Singapore  
Bachelor of Computing (Honours) in Computer Science | GPA: 4.9/5.0

Expected Graduation: 2027  
Singapore

EXPERIENCE

- TikTok  
Machine Learning Engineer Intern (Search)

Jan 2026 — May 2026  
Singapore

  - Incoming intern under the Vertical Search Team (Photo/Text & Live Search).
- Nika (Geospatial SaaS Startup)  
Software Engineer Intern

May 2025 — Aug 2025  
Singapore

  - Developed a multimodal geospatial AI agent that performs automated geospatial data analysis (inspired by Cursor AI) using different types of data formats, powering **100+ benchmark tasks with 92% workflow accuracy**.
  - Engineered a two-stage **dataset search and retrieval system** using a fine-tuned bi-encoder retriever and cross-encoder reranker, achieving **90% precision@5** within 1,000 ms end-to-end latency.
  - Containerised inference pipeline** (Docker & Kubernetes) and instrumented Prometheus metrics and Grafana dashboards for real-time monitoring of throughput, memory usage and latency distributions.
  - Developed an automated evaluation suite to measure precision@k, P99 latency and trigger regression gates for output drift.
  - Developed a managed DuckDB store, featuring persistent WebSocket connections and a containerised FastAPI service.
  - Created an **AI chatbot** with a multi-agent orchestration pipeline that translates natural language queries into SQL.
- Source Academy (NUS)  
Machine Learning Engineer (Part-time)

Jan 2025 — May 2025  
Singapore

  - Developed an AI tool to provide **automated feedback** on students' code submissions in Source, a JavaScript variant.
  - Built full-stack functionality, including backend logic for comment generation, database storage, and a frontend interface for teaching assistants to review AI-generated feedback.
  - Performed prompt engineering and **fine-tuned LLaMA 2 model using LoRA** to reduce token usage by providing additional Source language context, cutting AI processing costs.
- GovTech Singapore  
Software Engineer Intern (TechPass)

Apr 2024 — Aug 2024  
Singapore

  - Automated the creation of test accounts in bulk, reducing manual support tickets by **75%** across 10+ agencies.
  - Provisioned infrastructure as code** with Terraform and deployed microservices on AWS ECS.
  - Built Angular components with real-time updates, dynamic search filters, and pagination to simplify account management.

PROJECTS

- LLM Inference Engine | Python, gRPC, ONNX Runtime, Kubernetes

  - Built a gRPC inference API with dynamic batching and ThreadPoolExecutor, achieving sub-200ms latency.
  - Implemented ONNX Runtime optimisation with dynamic quantisation, beam search decoding, and KV-caching.
  - Deployed on Kubernetes with rolling updates, horizontal auto-scaling and health probes.
- Learny LMS | Node.js, Next.js, PostgreSQL, Express

  - Developed and launched an LMS for freelance tutors with course management, assignment tracking, payments tracking, and student progress monitoring, achieving adoption by 50+ users.

TECHNICAL SKILLS

**Languages:** JavaScript, Python, Golang, Java, C++, Ruby, HTML/CSS  
**Technologies/Frameworks:** React.js, Node.js, GraphQL, Spring Boot, FastAPI, Django, Flask, Docker, Kubernetes  
**Developer Tools:** Git, VS Code, Visual Studio, Firebase, Supabase, AWS DynamoDB, Terraform, Fireworks, Pinecone  
**Libraries:** NumPy, pandas, Matplotlib, PyTorch, LangChain, Vercel AI SDK