

Salman Naqvi | Computer Vision · NLP · Embodied AI Systems · Consultancy

+852 9586 7798 | Hong Kong SAR | s1148093@s.eduhk.hk | <https://snqv.dev/>

Building AI systems to see and understand the world better. Experience in embodied/spatial systems, distributed LLM training, and diffusion image generation. Has working Chinese fluency and provides a fast learning curve.
May 2027 expected graduation.

EDUCATION

The Education University of Hong Kong (EdUHK)

B.Sc. (Hons) in AI & Educational Technology

Hong Kong SAR

Sep 2023 – May 2027

- CGPA: 3.58/4.00 | Dean's List, EdUHK Full Entrance Scholarship
- fastai fellow, solveit alumnus

EXPERIENCE

VLM Researcher Intern

Focus Tech Lab

Jun 2025 – Aug 2025

HSITP (HK-SZ Innovation Technology Park), Hong Kong SAR

- In 2 months turned VLM spatial reasoning from black-box process to transparent debugging workflow
- With no model pretraining, established framework that improves VLM spatial reasoning at inference time
- Prevented estimated weeks of misdirected effort on diagnosing core failure point in robot arm pipeline

Embodied AI Researcher Part time

Multimodal Intelligence Lab, EdUHK

Oct 2024 – June 2025

EdUHK, Hong Kong SAR

- Improved spatial reasoning accuracy by 7% (27% → 34%) by finetuning Qwen2.5 VL-3B on RGB-depth images on an A100 GPU
- Derisked VLM spatial finetuning as a viable research direction

AI Engineer Intern

Keweya Education Technology Limited

Jul 2024 – Aug 2024

HKSTP (HK Science Technology Park), Hong Kong SAR

- Cut inference costs by 50% (HKD 14/hr vs HKD 28/hr) by implementing optimized GPU algorithms
- In 1 month, deployed a full-stack HTMX web app with custom LoRA modules, involving Tencent Hunyuan Photo Maker and Stable Diffusion on L40 and RTX 4090 GPUs

LLM Researcher Intern

Mobut Limited

Jun 2024 – Jul 2024

HKSTP, Hong Kong SAR

- Cut LLM training time by 80% (10hrs → 2hrs), saving HKD 1,200 in compute per run across 8xL4 GPUs, by replacing nested processing loops with vectorized operations
- Fine-tuned 70B+ parameter LLM in two hours at an ~75% lower compute cost by implementing FSDP+QDoRA distributed training

PROJECTS

HKD 20,000 Grant Recipient, EdUHK Sponsored Research Project

Nov 2025 – Jul 2026

- Assigned to produce a white paper on implementing AI systems for schools, businesses and educational institutions

GPU Algorithm Optimization

- Achieved 6x speedup (1.5s → 0.25s) on meanshift clustering by combining batch parallel distance computation, triangular kernel approximation, matrix multiplication over element-wise summation

DiffEdit: Text-based Image Editing

- Implemented DiffEdit paper from scratch, building a text-based diffusion image editing pipeline

SKILLS

- **ML/DL/AI:** PyTorch, fastai, Tensorflow, Hugging Face, NLP, OCR, Computer Vision
- **Agents:** Claude Code, OpenAI Codex, OpenClaw, SolveIt, Subagent orchestration, RAG, Agentic workflows
- **Infrastructure:** Linux, Unix, Bash, SSH, SQL, MongoDB, Git, Docker, CI/CD, GCP/Azure
- **Data & Tools:** Pandas, NumPy, Scikit-Learn, Gradio, Jupyter, HTML/CSS, nbdev
- **Languages:** English (native, IELTS 8.5/9.0), Mandarin (working proficiency, ~HSK5), Urdu (Fluent)