

Aayush Bajaj

📍 Sydney 📩 j@abaj.ai 📞 0481 910 408 📧 abaj.ai 🐧 github.com/abaj8494 💬 linkedin.com/in/abaj

Education

University of New South Wales
Master of Statistics

Sydney, Australia
Feb 2027 – Dec 2029

University of New South Wales
Bachelor of Computer Science (AI), Minor Mathematics

Sydney, Australia
Feb 2021 – Sept 2025

Skills

Deep Learning: PyTorch, TensorFlow, HuggingFace Transformers, nnU-Net, U-Net, CNNs, LLMs, RAG

ML Engineering: Python, CUDA, HPC (H200/A200), Distributed Training, Model Optimisation, Hyperparameter Tuning

Computer Vision: Semantic Segmentation, 3D Medical Imaging, Data Augmentation, Preprocessing Pipelines

LLM & GenAI: LangChain, LangGraph, LlamaIndex, smolagents, OpenAI API, Prompt Engineering, Fine-tuning, RAG

MLOps & Cloud: Docker, AWS, Azure, Git, Experiment Tracking, Model Deployment

Projects

Biomedical Semantic Segmentation

KiTS19

Production-grade kidney and tumour segmentation from 3D CT scans on official KiTS19 leaderboard.

- Achieved 0.9129 Dice score (ranked #57 globally) by implementing nnU-Net with systematic hyperparameter optimisation on H200 GPUs.
- Reproduced state-of-the-art 3D medical imaging pipeline end-to-end; published technical report with ablation studies comparing 2D vs 3D architectures.

Full-Stack LLM Chat Application

Production-ready conversational AI with RAG pipeline, streaming responses, and persistent context.

- Architected full-stack chat application using TypeScript, Bun, Svelte, and Prisma ORM with PostgreSQL; integrated OpenAI GPT API with streaming.
- Implemented security middleware with rate limiting, response caching, and structured logging; validated with Zod schemas and comprehensive test suite.

Neural Networks: Zero to Hero

Karpathy

From-scratch implementations of neural networks culminating in a GPT language model.

- Built micrograd (autograd engine), makemore (character-level LM), and nanoGPT following Andrej Karpathy's curriculum.
- Implemented backpropagation, attention mechanisms, and transformer architecture from first principles in pure Python/PyTorch.

Dead Tree Segmentation

Kaggle

Semantic segmentation of live/dead trees from satellite imagery with severe class imbalance.

- Achieved 0.71 IoU on 97% imbalanced dataset by implementing U-Net from scratch (31M parameters) with Dice + Focal loss combination.
- Parallelised training pipeline on CUDA-enabled HPC; applied morphological post-processing for boundary refinement.

Augmenting Infrastructure

Monorepo

Technical knowledge platform with ML/AI research notes, graduate-level mathematics, tutorials, and paper summaries.

- Architected full-stack content platform with Hugo static generation, Emacs org-mode authoring, and bidirectional linking across 1,073+ pages.
- Built literate programming environment with org-babel for executable Python/Julia notebooks, replicating Jupyter workflows with version control.

Certificates

AWS Solutions Architect (SAA-C03)
Amazon Web Services

Jan 2026 – Feb 2026

AI Agents Course
Hugging Face

Jan 2026 – Jan 2026