

Aryan Kekre

I'm an AI engineer with hands-on experience in AI training, RLHF, data annotation, model evaluation, prompt engineering, and fine tuning LLMs

LLM training, RLHF, and RAG for production · \$4/hr · Available

Kolkata · github.com/aryank-02

SKILLS

Python, PyTorch, RAG, Prompt Engineering, RLHF, Data Annotation

ABOUT

I'm an AI engineer with hands-on experience in AI training, RLHF, data annotation, model evaluation, prompt engineering, and fine tuning LLMs.

PROJECTS

Document Verifier SLM (Deep Learning / Transformer Architecture)

PyTorch · HuggingFace

Built a production-style identity document verification system using a custom Dual-Encoder Cross-Attention Transformer trained from scratch in PyTorch on 500K synthetic samples.

Designed a character-aware tokenizer and transformer architecture achieving 93% test accuracy for Aadhaar/document forgery detection. Implemented curriculum learning, EMA, focal loss, and hard negative mining; developed a three-tier inference pipeline with explainable field-level attribution.

Neural Archive v3.0 (RAG & LLM Infrastructure)

FastAPI · ChromaDB · Gemini API

Engineered an end-to-end RAG platform using FastAPI, ChromaDB, Gemini embeddings, and IBM Docling for OCR-aware ingestion of technical documents. Implemented semantic chunking, vector persistence, neighboring-context retrieval, and 3D embedding visualization via PCA + Plotly.