

ANKIT AGLAWE

Manchester, UK | +44 7570 293376 | aglawe.ankit@gmail.com | ankitaglawe.com
linkedin.com/in/ankit-aglawe | github.com/ankit-aglawe | huggingface.co/AnkitAI

SUMMARY

AI engineer with 6 years building and shipping LLM-powered production systems. Own 9 microservices at Verofy that handle fraud scoring, document extraction, compliance screening, and conversational AI, all routing across OpenAI, Claude, and Grok. Open source maintainer with 26,000+ PyPI downloads and 4 published Hugging Face models. Side project: an autonomous FTSE 350 trading agent on OpenClaw. MSc in AI, First Class Distinctions, Royal Holloway.

TECHNICAL SKILLS

Languages: Python, SQL, TypeScript

ML / AI: LLMs, RAG, Agentic AI, Prompt Engineering, NLP, Computer Vision, Deep Learning, Generative AI, Sentiment Analysis, Fraud Detection, Explainable AI (SHAP)

LLM and agents: OpenAI API, Anthropic Claude API, xAI Grok, MCP, Agno, LangChain, LlamaIndex, Hugging Face Transformers, ONNX Runtime, Structured Output, Function Calling

Frameworks: FastAPI, PyTorch, TensorFlow, scikit-learn, CatBoost, XGBoost, LightGBM, Optuna, Polars, Pandas, Streamlit

Cloud and data: AWS (S3, ECS, CodeDeploy, CloudWatch), Docker, CI/CD, Redis, MySQL, Qdrant, PostgreSQL, MongoDB, Kafka, Sentry

EXPERIENCE

Verofy Ltd

Manchester, UK

Senior AI Engineer

Feb 2024 – Present

- Built a multi-tenant RAG chatbot for merchant support (Agno, Qdrant, hybrid retrieval with BM25 reranking). Support response time dropped 40%, answer accuracy up 35%.
- Wrote the fraud detection engine: CatBoost/XGBoost ensembles with SHAP explainability and Optuna tuning. Scores transactions in real time, sends automated daily risk reports via email and Slack.
- Created the LLM orchestration layer used by all 9 microservices. Routes between OpenAI, Claude, and Grok with automatic fallback on provider failure.
- Designed the adverse media screening pipeline. HuggingFace sentiment models pre-filter search results before GPT-4o runs entity analysis against Companies House records. Pre-filtering alone cut LLM API spend 50% and halved manual review time.
- Wrote a statement extraction system that reads PDFs from 9 UK payment processors (Worldpay, Dojo, Barclaycard, Elavon, Clover) using Claude Sonnet, maps results to 400+ internal card categories, and generates pricing recommendations.
- Own the merchant monitoring system. Scores risk across 7 dimensions with 40+ configurable rules, processed on Polars for fast daily batch runs.
- Developed a review aggregation service: scrapes Google, Trustpilot, and TripAdvisor, runs ONNX-optimised sentiment and zero-shot classification through async Redis queues.
- Automated proof-of-address verification. Playwright screenshots UK directories, Tesseract extracts text, fuzzy matching validates addresses. Circuit breakers handle provider outages.
- Shipped an ad generator that turns merchant profiles into marketing visuals via Stability AI SD3.5 or DALL-E 3, with LLM-generated prompts keyed to merchant category codes. Design costs down 60%.

Asharth Asia PTE LTD

Singapore

Data Scientist

Aug 2021 – Aug 2022

- Worked on the autonomous vehicle stack: YOLOv4 for object detection (92% accuracy) and a Transformer-based OCR system for traffic sign recognition (98%).
- Improved food delivery routing with K-means clustering and geospatial analysis. Delivery times dropped 10%, efficiency up 25%.
- Shipped TensorFlow Lite models for on-device inference in the mobile app. A/B-tested personalised notifications; engagement up 20%, retention up 30%.

Shyena Tech Yarns PVT LTD

Pune, India

Data Scientist

Sep 2019 – Aug 2021

- Built a transaction risk scoring system (XGBoost, CatBoost, LightGBM ensembles) across 11.5M transactions. 95% accuracy, cut fraud losses 40%.
- Created anomaly detection for industrial vibration data using LSTMs and autoencoders. False positives dropped 15%.
- Replaced standard OCR with TableNet and Faster R-CNN for financial document extraction. Precision up 20%.

PROJECTS

OpenClaw LSE Trading Agent

Autonomous trading agent built on the OpenClaw framework that screens FTSE 350 stocks on the London Stock Exchange. Pulls real-time and historical market data from Yahoo Finance, computes technical indicators (Bollinger Bands, RSI, MACD, 50/200-day EMA crossovers), then scores financial news sentiment with an LLM before confirming trade signals. Includes position sizing rules, automatic stop-losses, and a backtesting harness against 5 years of historical LSE data. Published as a skill on ClawHub.

OPEN SOURCE

Python packages (26,000+ combined PyPI downloads)

- **emotionclassifier** (8,400+): Classifies text into 6 emotions using fine-tuned DeBERTa. CLI, batch processing, fine-tuning support.
- **configconverter** (6,600+): Converts between JSON, YAML, TOML, INI, and XML. CLI and Python API.
- **textpredict** (5,500+): NLP toolkit for sentiment, emotion, NER, and zero-shot classification with model comparison.
- **sentimentpredictor** (5,300+): Sentiment analysis with ONNX-optimised RoBERTa for fast inference.

Hugging Face

- 4 published transformer models including DeBERTa-XLarge (800M params) emotion classifier and DistilBERT financial news sentiment model.
- CategorizedTextReviews dataset (7,000+ downloads).

EDUCATION

University of London, Royal Holloway | Egham, UK

Sep 2023

MSc Artificial Intelligence – First Class Distinctions

CERTIFICATIONS

- AWS Certified Solutions Architect – Associate
- Introduction to Generative AI | Introduction to Large Language Models (Google)