

AAKARSHAN BHATIA

AI Engineer | LLMs, RAG and Machine Learning

London, United Kingdom | aakarshan1995@gmail.com
[linkedin.com/in/aakarshan-bhatia](https://www.linkedin.com/in/aakarshan-bhatia) | github.com/aakarshan1995

PROFESSIONAL PROFILE

AI Engineer with an MSc in Artificial Intelligence (Distinction) and 5+ years of technology experience across AI engineering and fintech quality assurance. Founder of Adaptive Labs, designing and building Large Language Model (LLM) applications using Retrieval-Augmented Generation (RAG), AI agents and tool calling, with full-stack delivery across APIs, web and mobile. Brings a production-readiness mindset from testing business-critical payment systems at FIS Worldpay.

TECHNICAL SKILLS

Languages: Python, SQL, JavaScript, TypeScript, HTML, CSS

Generative AI: Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), AI agents, multi-agent workflows, tool/function calling, Model Context Protocol (MCP), prompt engineering, embeddings, vector search, semantic search

LLM Platforms and Search: OpenAI, Claude, Gemini, OpenRouter, Ollama, Weaviate, Elasticsearch

Machine Learning and Data Science: scikit-learn, pandas, NumPy, Natural Language Processing (NLP), predictive modelling, classification, regression, clustering, anomaly detection, feature engineering, model evaluation, hyperparameter tuning, statistical modelling, TF-IDF, logistic regression, probability calibration, Jupyter, Tableau

Backend, Web and Mobile: FastAPI, Flask, Node.js, REST APIs, React, Flutter, PostgreSQL, Firebase

Cloud and DevOps: AWS (S3), Docker, Git, cloud and on-premise deployment

Testing and Delivery: Selenium, test automation, operational acceptance testing (OAT), risk-based testing, Agile, SAgE

PROFESSIONAL EXPERIENCE

Founder and AI Engineer | Adaptive Labs | London, United Kingdom (Hybrid) Sep 2025 - Present

- Founded a software and AI technology firm, leading design and delivery of production-ready AI applications.
- Design and build LLM systems using Retrieval-Augmented Generation (RAG), AI agents, tool/function calling and agentic orchestration, covering embeddings, semantic search, prompt engineering and model evaluation.
- Integrate LLMs with external APIs, databases, vector stores and business systems, and prototype agent-to-tool connectivity with the Model Context Protocol (MCP).
- Develop APIs and full-stack web and mobile applications, taking AI products from technical design to deployment.

OAT Test Analyst (Quality Assurance) | FIS Worldpay | London, United Kingdom 2020 - 2024

- Led operational acceptance testing (OAT) across 30+ enterprise fintech projects for critical payment infrastructure.
- Improved testing efficiency by 30% by introducing Selenium test automation and streamlining test processes.
- Managed an enterprise fintech project end to end, introducing new operating methods that improved efficiency 4x.
- Designed risk-based test strategies for cloud and on-premise systems across test, pre-production and production.
- Partnered with architects, engineering teams and business stakeholders to validate system behaviour, resilience and deployment readiness, catching critical defects before release.
- Mentored new analysts and supported Agile and SAgE delivery practices.

PROJECTS

LectureMate: RAG-Based AI Study Assistant | Personal Project

- Built a RAG-based study assistant with FastAPI and React/TypeScript, using embeddings and vector search to retrieve relevant lecture content and generate grounded responses.
- Integrated Weaviate, Ollama/OpenAI, document chunking and AWS S3; designed for containerised cloud deployment.
GitHub: github.com/aakarshan1995/Lecturemate_RAG_based_project

MSc Dissertation: Mental Health Risk Assessment from Social Media Text | Queen Mary University of London

- Designed a dual-stage (coarse-to-fine) NLP classification pipeline using TF-IDF features and calibrated logistic regression to triage posts across imbalanced mental health classes.
- Improved macro-F1 by 0.075 over a single-stage baseline (paired bootstrap, $p < 0.01$) and raised recall at 90% precision from 0.50 to 0.62; assessed ethical risks including stigmatisation.

Hybrid Recipe Search Engine | Personal Project

- Built a hybrid retrieval system combining Elasticsearch BM25 with LSA semantic search (TF-IDF, TruncatedSVD).
- Evaluated retrieval using precision, recall, F1 and MAP; hybrid search achieved 0.78 recall on the project evaluation set.
GitHub: github.com/aakarshan1995/hybrid-recipe-search-engine

EDUCATION

MSc Artificial Intelligence (Natural Language Processing), Distinction

2024 - 2025

Queen Mary University of London

BSc Computer Science, First-Class Honours | University of Hertfordshire

CERTIFICATIONS AND TRAINING

Software Engineering Virtual Experience Programme, Goldman Sachs (2024)