

Luv Agrawal

London, UK | luvga07@gmail.com | +44 7823 705339 | linkedin.com/in/luvagrawaluwl | github.com/luv-codes

SUMMARY

AI and Data Engineering intern who shipped two production systems in parallel: a GraphRAG platform over 1,000+ NHS clinical documents (query accuracy improved 74%) and a contract intelligence product launched with a UK Tier 1 bank as its first client, including fine-tuning a small Qwen3-4B model. Final-year BSc Computer Science student at the University of West London with a 90%+ average and applied experience across retrieval systems, knowledge graphs, model fine-tuning, and event-driven AWS pipelines. Seeking a 2027 graduate role in AI Engineering, Data Engineering, or technology consulting.

WORK EXPERIENCE

AI and Data Engineering Intern, HaloRFP, London, UK

July 2026 - September 2026

- Built a contract intelligence platform for commercial real estate and construction from initial design to first client deployment, with a UK Tier 1 bank (name under NDA) as the launch client for contract, lease, and tender analysis.
- Fine-tuned a small Qwen3-4B language model as part of the contract intelligence work.
- Engineered document ingestion for contracts, leases, and tender packs: PDF and DOCX parsing, clause-aware chunking of legal text, and structured metadata capture (parties, key dates, document type).
- Implemented LLM-based extraction converting contract text into a defined JSON schema of obligations, break clauses, payment terms, and risk flags for downstream legal review.
- Designed the retrieval and question answering layer using semantic search with clause-level metadata filtering, returning every answer cited back to its exact source passage.
- Authored the full technical handover: architecture, runbook, troubleshooting guide, API reference, decisions log and cost model, enabling the team to continue the build without loss of context.
- Delivered in Agile Scrum sprints, with sprint planning and task tracking managed in ClickUp.

AI and Data Engineering Intern, I/O Atelier, London, UK

June 2026 - September 2026

- Improved clinical query accuracy by 74% (94/100 vs 54/100 correct at k=5) by building a GraphRAG layer over a BM25 and vector search Hybrid RAG baseline, using DeepSeek-V3 entity extraction and a NetworkX knowledge graph persisted in pgvector (23,576 nodes, 57,050 edges).
- Raised retrieval confidence by 27 percentage points (81% vs 54%) with a metadata and entity tagging pipeline enabling graph traversal across 8-14 documents per query; evaluated 500 clinical queries across five top-k settings.
- Eliminated three categories of GraphRAG false positives (minimum keyword overlap threshold, confidence floor, citation breadth recalibration) before scaling the system to 1,000+ NHS documents.
- Built event-driven AWS ingestion pipelines (S3, Lambda, IAM, CloudWatch, Postgres RDS) processing GeoJSON datasets across 20+ field types with coordinate validation and least-privilege IAM policies.
- Automated conversion of raw estate data (buildings, paths, cost plans, condition reports) into structured node and relationship JSON through Neo4j ingestion logic.

PROJECTS

GraphRAG MCP Server (internship project, private company repository)

2026

Python, FastAPI, MCP (Model Context Protocol)

- Wrapped the production retrieval engine as an MCP server, exposing graph traversal and document search as callable tools for AI agents and assistant integrations.
- Defined tool schemas and request handling on top of the FastAPI service, letting an AI assistant answer questions over the knowledge graph without direct database access.

AWS Serverless ETL Pipeline - Synthetic Patient Data

2026

Python, Amazon S3, AWS Lambda, IAM, CloudWatch, Boto3, Pandas

- Processed 100,500 synthetic patient records in under 2 seconds with an S3-triggered Lambda function; produced age band distributions, 30-day readmission rates, and a 26-field data quality report.

EDUCATION

BSc (Hons) Computer Science, University of West London

September 2024 - June 2027 (Expected)

Final year. Current average: 90%+. Key results: Theory of Computation 100%, Artificial Intelligence 92%, Mathematics 91%, Algorithms and Data Structures 90%, Programming 89%.

SKILLS

AI and RAG: GraphRAG, Hybrid RAG, Retrieval Augmented Generation, BM25, Vector Search, Semantic Search, Embeddings, LangChain, RAG Evaluation, Entity Extraction, Knowledge Graphs, NetworkX, LLM Fine-Tuning, Qwen3, DeepSeek-V3, MCP (Model Context Protocol), Prompt Engineering

Programming: Python, SQL, FastAPI, REST APIs, Pandas, NumPy, Boto3, GeoPandas, Matplotlib, Seaborn

Cloud: AWS, EC2, Terraform, Lambda, S3, IAM, CloudWatch, Serverless Architecture, Event-Driven Pipelines, Docker

Databases: PostgreSQL, Database Sharding, Supabase, pgvector, Neo4j

Tools: Git, GitHub, VS Code, Jupyter, ClickUp, Agile Scrum, Technical Documentation

ACHIEVEMENTS AND CERTIFICATIONS

- BMW Group UK Data Analytics Final Assessment Centre (December 2025): reached the final stage (top 1% of applicants nationally) and presented AWS architecture solutions to senior assessors.
- AWS Cloud Practitioner Foundations: in progress (2026).
- Mathematics Olympiad: Distinction (Grades 5 and 7). Microsoft x EDI Diversity in Tech Competition participant (2025).