ROHAN VYAS

rohanvyas@hotmail.com | linkedin.com/in/rohanvyas0 | rohan-vyas.com

EDUCATION

University of Nottingham, BSc Hons Computer Science (2:1) *Sep 2022 – Jul 2025*

- Relevant modules: Algorithms, Data Structures and Efficiency; Fundamentals of Artificial Intelligence; Designing Intelligent Agents; Computer Graphics (C++)

TECHNICAL SKILLS

Languages: Python, TypeScript, Go, C, C++, SQL, Scala
ML & AI: LLM integration (Gemini, GPT), LLM agents and orchestration, Google ADK, model evaluation, labelled data curation, schema-constrained output, RAG and embeddings, scikit-learn, NLP (TF-IDF, cosine similarity)
Data & Systems: Spark, Databricks, Delta Lake, PostgreSQL, pgvector, Redis, FastAPI, Next.js, Docker, AWS, Playwright, Git, Linux

EXPERIENCE

Apple | AI Security Engineer *London | Jan 2026 – Present*

- Evaluate autonomous LLM agents that reproduce externally reported security findings, auditing their verdicts against execution traces under a disposition taxonomy
- Triage high-volume security detections under a multi-class labelling scheme, recording the evidence behind each verdict, then automated that review in a Python pipeline feeding a metrics dashboard
- Overturned the highest-volume finding class by empirically retesting whether flagged tokens were live, replacing prefix matching with a validity check and removing a systematic false-positive source
- Built the retrieval layer for a batch LLM classifier of incoming security reports, embedding the CWE catalogue and constraining output to a validated taxonomy
- Build autonomous agents on Google ADK and Model Context Protocol, orchestrating Playwright browser automation to reproduce and validate findings end to end

Arvo (construction contract intelligence) | Co-Founder & AI Engineer *London | Apr 2026 – Present*

- Built a contract extraction pipeline turning NEC4, JCT and FIDIC PDFs into typed clause records of obligations, deadlines, payment terms and risk, validating every model-supplied date and grade before storage
- Shipped a six-class event classifier whose confidence sets severity and whose category gates routing, constraining the natural-language query layer to validated filters so every figure shown is computed deterministically

Kaizo (AI operations for dental practices) | Co-Founder & CTO *London | Sep 2024 – Present*

- Root-caused a false alarm of 29,034 uninvoiced treatments to a join key the data sync never populates, rewriting the detector as 12-month per-visit matching that returned a 5,219-visit review list from 14,089
- Shipped safety detection over generated text: six clinical policy categories returned with the offending excerpt and the send blocked, plus eight-category prompt-injection detection clamping agent autonomy
- Risk-classified 41 agent actions into three tiers, routing high-risk ones into a role-gated approval queue, and scored the two-intent routing rule against a 53-case labelled suite

Apple | Information Security Engineer Intern *London | Jun 2025 – Sep 2025*

- Engineered a parameterised Spark and Delta Lake pipeline (Scala, Python) that archived analyst queries and re-ran them on schedule, automating a manual re-run measured at 10 to 20 minutes per query
- Cut false positives on scheduled detection queries by tuning benign-pattern and benign-domain exclusions against live hits, then standardised output onto the partner team's shared emitter template

Accenture (6point6) | Data Science & Cyber Security Intern *London | Jun 2024 – Aug 2024*

- Benchmarked five generative models including Gemini and GPT-4 for gender bias, running a fixed 40-prompt suite identically per model and scoring ~2,000 generated text units on two automated measures
- Extended an inherited bias-scoring harness with a Gemini client replacing the retired Bard integration, so all five models scored through one comparable pipeline, and documented the study's eight stated limitations

PROJECTS

Conversational Intent Classifier | Python, scikit-learn, NLP *Oct 2024 – Dec 2024*

- Built a TF-IDF and cosine-similarity intent classifier across 17 classes, requiring 0.4 similarity to commit to an intent and routing anything below it to retrieval over a 1,473-row question-answer corpus

Multi-Level Feedback Queue Scheduler | C, POSIX Threads, Data Structures *Oct 2023 – Jan 2024*

- Built a process scheduler in C over ten versions: 16 priority levels with periodic boosting, round-robin time slicing, per-device I/O queues, and POSIX threads synchronised by mutexes

Gaffer, Autonomous Fantasy Football Manager | Python, Forecasting, Calibration *Aug 2026 – Present*

- Modelled player output with sample-size-weighted shrinkage towards fitted position priors, and built a calibration harness measuring predicted probabilities against observed outcomes, discounting buckets under 30 samples

LEADERSHIP

Nottingham Asian Society | Vice President *Apr 2023 – Jul 2024*

- Led a 20-member committee serving 500+ students; organised a £30k formal ball that netted £5k profit through vendor negotiation and ticket strategy