

Profile

AI/ML engineer with three years' experience taking machine-learning and LLM systems from ambiguous enterprise requirements through production deployment. Delivered yieldHUB's first production ML capability and production computer-vision systems, specialising in evaluation, failure handling, and human-in-the-loop design. MSc Artificial Intelligence, University of Edinburgh.
Seeking: Part-time AI/ML opportunities during MSc; full-time from Sept 2027 | **UK work authorisation:** unrestricted (Irish citizen)

Experience

Data Scientist - AI/ML Systems Engineering <i>yieldHUB</i>	2023 – 2026 <i>Limerick, Ireland</i>
- Owned a production X-ray microchip inspection system end-to-end, from automated ingestion to user-facing interface, reducing manual review effort by approximately 80%. - Delivered yieldHUB's first production ML capability: large-scale anomaly detection used daily by engineering teams and supporting automotive quality workflows. - Redesigned a production ingestion pipeline after customer-reported performance issues, improving throughput by approximately 16× and reducing compute costs by approximately 60%. - Designed and prototyped a modular MCP-based agent architecture, defining reusable LLM abstractions aligned with existing APIs and platform constraints. - Developed ML systems directly with enterprise customers, refining ambiguous requirements and challenging infeasible approaches; represented yieldHUB at industry events and delivered technical presentations on AI in manufacturing. - Designed human-in-the-loop workflows with explicit failure handling and debugged time-critical production issues to maintain reliability in high-stakes manufacturing environments. - Mentored interns and supported internal adoption of modern ML and GenAI practices.	
Data Engineering Intern <i>Groupon</i>	Mar 2022 – Aug 2022 <i>Remote</i>
Built self-service GCP data-transfer tooling for non-technical users and automated reliable, repeatable backups with Airflow scheduled jobs.	

Selected Projects

Auditable Agent Loop: Adaptive investigation agent separating LLM planning/interpretation from deterministic computation, with evidence-level provenance, contradiction handling, replayable investigations, typed termination, and evaluation over SEC EDGAR and tabular data.

NanoGPT From Scratch | Live: Reproduced GPT-2 124M from scratch (3.0503 validation loss; 0.3043 HellaSwag); paired-seed ablations found optimiser choices dominated architecture (WSD: -0.1034 validation loss vs baseline), with 95.1% scaling efficiency across 8 GPUs.

RAG Eval & Observability | Live: Deployable RAG platform with multiple retrievers, persisted offline evaluation, regression comparison, query tracing, metrics, CI evaluation gates, and end-to-end testing.

Stepwise: Multimodal RAG using transcript/frame ingestion, HyDE retrieval, cross-encoder re-ranking, and cited timestamped visual evidence.

Technical Skills

Languages: Python, TypeScript, SQL **LLM Systems:** RAG, agentic workflows, MCP, evaluation, guardrails, structured outputs

Engineering: FastAPI, Next.js, REST APIs, Docker, Airflow, CI/CD, observability, production deployment

ML/Data: PyTorch, Transformers, PostgreSQL, pgvector, Qdrant, anomaly detection, model evaluation, semantic/hybrid search, re-ranking

Education

MSc Artificial Intelligence <i>University of Edinburgh</i>	2026 – 2027 <i>Edinburgh, UK</i>
Current coursework in Probabilistic Machine Learning, Machine Learning Systems, and Accelerated Natural Language Processing; planned study includes Advanced NLP, Robotics, and Reinforcement Learning, followed by an MSc dissertation.	
BSc Data Science & AI (2:1) <i>University College Cork</i>	2019 – 2023 <i>Cork, Ireland</i>
Foundation in statistical inference, machine learning, feature engineering, model evaluation, probabilistic modelling, and applied work in Python, SQL, and R.	