

# Joel Smith

## SUMMARY

Data Science & Artificial Intelligence MSc graduate with a strong focus on building reliable, well-structured machine learning systems. Experienced in translating mathematical models into clean, reproducible code, and in designing end-to-end pipelines spanning data, training, evaluation, and deployment. Demonstrated depth across generative modelling and applied ML systems, with hands-on experience using PyTorch and modern ML tooling.

## Education

<b>University of Liverpool</b> <i>MSc in Data Science &amp; Artificial Intelligence, Distinction</i>	Sept 2024 - Sept 2025
<b>University of Chester</b> <i>BSc in Economics, First Class</i>	Sept 2020 - July 2023

- Awarded the University Prize for Economics

## Projects

<b>LLM-Based Financial Advisor Web App</b>	GitHub Repo
<ul style="list-style-type: none"><li>• Designed and implemented a full-stack LLM-based financial Q&amp;A system to reduce information asymmetry in retail finance, integrating retrieval-augmented generation with real-time market APIs and vector search (ChromaDB/pgvector) to ground responses in up-to-date, verifiable data.</li><li>• Engineered a production-oriented inference pipeline with a FastAPI backend and React frontend, containerised via Docker and deployed on a local Kubernetes (kind) cluster to support modular services and reproducible serving.</li></ul>	
<b>Manipulation Lab</b>	GitHub Repo
<ul style="list-style-type: none"><li>• Designed and implemented a modular benchmarking framework for robotic imitation learning in NVIDIA Isaac Sim, supporting end-to-end pipelines for data collection, training, and evaluation.</li><li>• Reduced experimentation overhead by shifting task and training configuration from code to Hydra-based configs, enabling rapid iteration and reproducible comparisons across benchmarks.</li><li>• Improved maintainability by decoupling environment, dataset, and model components, reducing boilerplate and simplifying the addition of new tasks.</li></ul>	
<b>Generative Modelling from Scratch</b>	GitHub Repo
<ul style="list-style-type: none"><li>• Authored a blog-style series of deep technical studies on generative models (e.g., mixture models, VAEs, diffusion models) developing detailed mathematical derivations and explaining core probabilistic concepts from first principles.</li><li>• Produced supporting PyTorch/NumPy implementations and visualisations to illustrate the theory and compare learned generative distributions against real data.</li></ul>	
<b>Reinforcement Learning from Scratch</b>	GitHub Repo
<ul style="list-style-type: none"><li>• Authored a notebook-based series on reinforcement learning fundamentals, deriving and implementing core algorithms (bandits, dynamic programming, Monte Carlo methods) from first principles.</li></ul>	

## Machine Learning and Deep Learning Projects

University Coursework

- Implemented a range of applied machine learning and deep learning projects across multiple problem domains using classical and neural approaches.

## Experience

<b>Commercial Finance Analyst</b> <i>Matalan</i>	June 2024 – Sept 2024
	Liverpool, UK
<ul style="list-style-type: none"><li>• Queried, cleaned, and validated large-scale SQL datasets (10M+ rows) to support sales and promotional analysis, delivering time-sensitive insights to commercial finance stakeholders.</li></ul>	

**Team Supervisor** Sept 2020 – January 2024

*Costa Coffee*

Liverpool, UK

- Led a team in a high-throughput retail environment while completing a full-time undergraduate degree, developing management and communication skills.

## Skills

---

**Machine Learning:** PyTorch, scikit-learn, Torchvision, Hydra

**Data & Experimentation:** PostgreSQL, Google BigQuery, Microsoft SQL Server, Weights & Biases

**LLM & Retrieval:** OpenAI API, ChromaDB, pgvector

**Systems & Deployment:** FastAPI, Docker, Kubernetes (kind), Git, Linux, AWS (EC2)

**Languages:** Python, SQL