Mohamed Atoui

mohamedelamineatoui@gmail.com | +44 7985665024

GitHub | LinkedIn | Portfolio

SUMMARY

Machine Learning Engineer / Data Scientist with production experience building AI systems serving 500+ users. Skilled in developing ML applications, predictive modeling, and translating complex data into actionable business insights. Strong foundation in Python, SQL, PyTorch, and statistical analysis. Demonstrated ability to translate complex technical concepts for diverse audiences through teaching and client-facing product work. Proven ability to lead projects from research to deployment in fast-paced, client-facing environments.

EDUCATION

Royal Holloway University London - *3rd Year Computer Science & Artificial Intelligence BSc (Honors)*2023 - 2026

- **Grade:** First-Class Honors
- **Relevant Coursework:** Machine Learning, Deep Learning, Natural Language Processing, Computational Finance, Statistics

EXPERIENCE

Product Data Intern – Stats Perform Jun 2025 - Oct 2025

- Conducted market analysis using Monte Carlo simulations to model revenue potential for tech product distribution across cloud marketplaces (AWS), delivering insights to support go-to-market strategy.
- Analyzed large-scale sports datasets (NFL, NBA, MLB) using SQL and Python to identify statistical patterns, presenting findings to cross-functional teams to inform product development.
- Developed feature engineering workflows and implemented intelligent feature selection algorithms for sports analytics products.
- Performed data analysis to support SDK adoption strategy, translating technical capabilities into actionable product insights for business stakeholders.

Co-founder & ML Research Lead - *Nordlys Labs* Feb 2025 - Present

- Built and scaled AI infrastructure startup from concept to **500+ users**, validating product-market fit through iterative development and customer discovery; selected by Antler accelerator.
- Designed intelligent routing architecture that analyzes prompt complexity to dynamically allocate tasks across 30+ LLMs.
- Built production infrastructure with real-time monitoring dashboard enabling clients to track usage, costs, and performance metrics for data-driven business optimization.
- Engaged with clients to understand requirements and translate them into technical solutions, balancing business needs with engineering constraints.

Teaching Assistant – *Royal Holloway University of London* Oct 2025 - Present

- Mentored **50+** students in Object-Oriented Programming and Software Engineering, translating complex technical concepts into accessible learning materials.
- Provided structured feedback on technical assignments, demonstrating ability to communicate technical findings clearly and constructively.

Machine Learning Research – *ISG Smart Card & IoT Security Center* Dec 2024 - Sep 2025

- Researching adversarial robustness in ML-based intrusion detection for autonomous vehicles, analyzing failure modes under attack and distribution shift.
- Designed **CNN** and **LSTM** anomaly detection models; built evaluation framework for reliability testing on real-world automotive datasets.
- Contributing to academic reports and technical whitepapers, documenting analytical methodologies and communicating findings to technical stakeholders.

TECHNICAL PROJECTS

Vantage - AI ROI Calculator | Python, LangGraph, ChromaDB, Sentence Transformers | link

• Built an agentic AI assessment platform using LangGraph that orchestrates multi-step reasoning across feasibility scoring, three-year financial projections, and build-vs-buy analysis, compressing **40+** consultant hours into a **5-10 minute** evaluation.
• Implemented RAG pipeline with Sentence Transformers and ChromaDB to ground recommendations in **200+** real-world case studies, combining semantic retrieval with risk-adjusted financial modeling.

CrossGen - Crosss-Domain Idea Generator | *Python, FastAPI, Claude API, Pydantice* | link
• Built a 6-stage analogical reasoning pipeline that decomposes problems into domain-neutral relations and discovers solutions by mapping structural parallels across **42** curated domains using Gentner's Structure-Mapping Theory.
• Designed a hybrid deterministic + LLM architecture where TRIZ principles are looked up (not hallucinated) and Claude handles structural reasoning, with mandatory weakness documentation and falsifiable predictions for every generated solution.

KV-Cache Optimization Research | *Python, PyTorch, HuggingFace* | link
• Benchmarked four attention mechanisms (MHA, MQA, GQA, MLA) in iso-parameter 16M param transformers, Multi-Head Latent Attention achieved the best quality-efficiency trade-off, **13–26%** lower perplexity than MHA while reducing KV-cache by **72%**.
• Developing KV2State, a hybrid inference method that classifies attention heads as retrieval or streaming (via DuoAttention) and replaces streaming heads' growing KV cache with fixed-size O(1) recurrent state matrices, targeting Llama-3.1-8B with the goal of near-lossless quality at significantly reduced memory.

TECHNICAL SKILLS

• **Data & Analytics:** SQL, Python, Pandas, NumPy, Apache Spark (familiar), Hadoop (familiar)
• **ML/AI Frameworks:** PyTorch, TensorFlow, Scikit-learn, Hugging Face Transformers, LangChain, ONNX
• **ML Domains:** Predictive Modeling, NLP, LLMs, Computer Vision, CNNs, RNNs/LSTMs, Time Series Analysis
• **Visualization:** Matplotlib, Seaborn, Plotly, Dashboard Development
• **Infrastructure:** Docker, Git, GitLab CI/CD, Linux, FastAPI, AWS, Azure
• **Languages:** Python, Java, TypeScript, JavaScript, C, SQL
• **Languages Spoken:** English (Fluent), French (Fluent), Arabic (Native)