

AMIN LASRI

AI / ML Engineer

Ottawa, ON, Canada · aminlasri@gmail.com · linkedin.com/in/amin-lasri · github.com/217amin · aminanalyst.com

00 PROFILE

AI/ML Engineer specializing in production-grade retrieval pipelines and machine learning systems. Master's in Applied Econometrics combined with a Diploma in Artificial Intelligence (La Cité). Eight years as an Economics & Statistics educator, transitioning into AI engineering — bringing rigorous statistical inference to modern ML architectures.

Open to AI/ML Engineer, MLOps, and Applied Scientist roles.

01 PROJECTS

AeGis-RAG — Production-grade RAG System | Personal Project 2026

Hybrid retrieval (dense + BM25 with RRF fusion) + cross-encoder reranking + safe abstention to minimize hallucinations on policy QA. Three-stage pipeline allowing independent optimization of retrieval, ranking, and generation.

- **Recall@5 = 0.85, nDCG@5 = 0.80** — outperformed retrieval-only baseline on benchmarked policy QA dataset.
- **Latency = 0.73s with reranking** — exposed config flag to disable reranking for latency-sensitive deployments.
- **Skipped fine-tuning entirely** — proved hybrid retrieval (dense + BM25 + RRF) outperformed fine-tuned baselines on small structured datasets — saving training cost and infrastructure complexity.
- Stack: Python · LangChain · Chroma · SentenceTransformers · BM25 · Cross-Encoders · RAGAS · Ollama

Amazon ESCI Search Engine | Personal Project 2025

Fine-tuned Matryoshka embeddings on the Amazon ESCI e-commerce dataset to compress vector dimensions while preserving retrieval quality. Optimized inference cost at scale without sacrificing recall or nDCG.

- **91% dimension reduction** — while improving nDCG and Recall metrics on benchmarked test set.
- **Quantitative LLM evaluation** — designed evaluation harness using RAGAS, achieved Context Precision = 0.83.
- Stack: Python · sentence-transformers · PyTorch · Matryoshka embeddings

Housing Price Prediction | Personal Project 2025

End-to-end geospatial valuation pipeline using ensemble methods. Rigorous hypothesis testing and feature engineering across 80+ structural and geospatial variables.

- **R² = 0.9230 on held-out test set** — outperforming linear and tree-based baselines.
- **80+ engineered features** — combining structural, neighborhood, and geospatial clustering signals.
- Stack: Python · XGBoost · CatBoost · scikit-learn · geospatial clustering

Parking Vision System | Personal Project 2024

Deep learning vision pipeline for autonomous parking maneuvers, with real-time spatial obstacle detection.

- **88.33% accuracy**
 - Stack: Python · OpenCV · CNN architectures
-

02 EXPERIENCE

Economics & Statistics Educator | Ibn Zohr High School · Morocco 2016 — 2024

- Taught statistical inference, and quantitative methods to over 600 students across 25 cohorts.

- Designed curriculum that translated complex statistical theories — causal inference, hypothesis testing, regression — into actionable applied frameworks.
- Developed communication skills critical for translating ML systems into business value for non-technical stakeholders.

03 EDUCATION

Diploma in Artificial Intelligence | La Cité · Ottawa ON 2026

- Coursework: Machine Learning, Deep Learning, NLP, Computer Vision, MLOps, Production AI Systems.
- Capstone projects: AeGis-RAG, Amazon ESCI Search Engine, Parking Vision System.

Master's in Applied Econometrics | Hassan II University · Morocco 2011

- Specialization: time-series analysis, causal inference, panel data econometrics, statistical modeling.

04 SKILLS

LANGUAGES

Python · SQL · R

ML / DEEP LEARNING

PyTorch · scikit-learn · XGBoost · CatBoost · sentence-transformers

LLM / RETRIEVAL

LangChain · RAGAS · Chroma · BM25 · Cross-Encoders · Ollama · Matryoshka embeddings

COMPUTER VISION

OpenCV · CNN architectures

TOOLING

Git · Docker · Jupyter · pandas · NumPy

LANGUAGES (HUMAN)

French (native) · English (fluent) · Arabic (native)

Full project case studies available at aminanalyst.com