

Jaafar Benabderrazak

Senior Data Scientist/Machine Learning Engineer

Île-de-France, France

+33 6 59 98 45 38

✉ jaafarbenabderrazak.info@gmail.com

🌐 medium.com/@jaafar-benabderrazak

in linkedin.com/in/jaafar-benabderrazak

🐙 github.com/jaafar-benabderrazak



Profile

With 6 years of experience in Machine Learning and Data Engineering, I specialize in developing AI-powered solutions across finance, energy, and transportation sectors. Expert in AWS cloud technologies and LLM integration, with proven ability to reduce fraud rates and optimize operational efficiency. Strong background in MLOps implementation and predictive analytics, combining technical expertise with practical business impact. Holds multiple AWS certifications and has authored ML-focused articles with over 10,000 reads.

Technical Skills

- **Programming & Data Science:** Python, TensorFlow, Scikit-Learn, XGBoost, Random Forest, LSTM, NLP (spaCy, gensim, word2vec, topic modeling), LangChain, RAG Framework
- **MLOps & Workflow Orchestration:** AWS SageMaker, AWS Lambda, AWS Bedrock, AWS CDK, Terraform, MLflow, DVC, ZenML, GitLab CI/CD, Databricks
- **Cloud & Big Data:** AWS (S3, EC2, SageMaker, Bedrock), PySpark, Hadoop, Databricks, Vector Databases, OpenSearch
- **Databases & Querying:** SQL, Text-to-SQL systems, Neo4j (CypherQL), Knowledge Base Integration
- **Data Engineering & Analytics:** Data pipelines, ETL/ELT, Data Quality, Fraud Detection Systems, Predictive Modeling
- **Web & Automation:** Streamlit, Flask, Selenium, API Development, Web Scraping, Document Processing Automation

Professional Experience

Mar 2025 – **IA/LLM Engineer**, *Stellantis Financial Services (via Reply)*, Paris, France

Sep 2025 **Context:** SFS is the financial arm of Stellantis, providing leasing, insurance, and credit services. Mission carried out within a team of 2 (MLE, product owner).

- Architected and developed a Text-to-SQL system enabling 40+ financial analysts to query complex financial databases in natural language, reducing dependency on IT teams
 - Integrated semantic parsing and query optimization, achieving query accuracy of 85.6% and reducing average query response time from minutes to under 20s
 - Built an evaluation framework with metrics on accuracy, latency, and user satisfaction; decision-making cycle reduced by 27% thanks to faster insights
 - Delivered production-ready pipelines on AWS, ensuring scalability and compliance with financial standards
- Stack: Python, AWS Lambda, LangChain, OpenSearch, RAG Framework, Fine-tuned LLMs, AWS Bedrock*

Feb 2025 – **Machine Learning Engineer**, *Stellantis (via Reply)*, Paris, France

Jun 2025 **Context:** Mission within the R&D division (team of 3 including ML engineer, data engineer, and front-end engineer) focused on predictive maintenance and anomaly diagnostics for connected vehicles.

- Designed and deployed an LLM-powered diagnostic system analyzing vehicle sensor streams to detect and explain anomalies in natural language
- Implemented a knowledge retrieval pipeline linking anomalies with troubleshooting docs and repair procedures
- Added root cause analysis layer translating sensor data into actionable repair insights
- Achieved 72.8% anomaly detection accuracy, improving reliability of diagnostics and technician efficiency
- Reduced diagnostic time by 24.5% and cut misdiagnosis rate by 14.2%, saving estimated €1.1M annually in warranty claims

Stack: Python, AWS SageMaker, AWS Bedrock, Vector Databases, RAG Framework, AWS CDK

Sep 2024 – **Machine Learning Engineer**, *SNCF (via Talan)*, France

Dec 2024 **Context:** Mission within the TER division, working in a data team of 3 (1 DS, 1 DE, 1 PO) to combat ticketing fraud representing multimillion-euro losses annually.

- Reduced fraud rate from 7.2% to 5.9% through hypothesis testing and optimized controls
- Developed predictive models (RF, XGBoost, LSTM) with exogenous variables, achieving AUC-ROC > 0.87
- Regression models predicted fraud volumes with RMSE of 14.7%, enabling better inspector resource allocation

Stack: Random Forest, XGBoost, LSTM, Python

Jun 2023 – **Machine Learning Engineer**, *GRT Gaz (via Talan)*, Paris, France

Aug 2024 **Context:** Mission in R&D innovation program within a 6-person team (DS, MLE, software engineers, PM), aiming to enhance operational safety of France's gas transport network.

- **Gas Propagation Tool:** Developed predictive models for pipeline gas dynamics (time series forecasting), improving anticipation of incidents
- **LLM Chatbot:** Led development of a Confluence-integrated assistant, reducing support response times by 38.5%
- **Support & Maintenance:** Continuous monitoring and retraining of prediction tools; introduced TDD practices, improving maintainability

Stack: Python, AWS SageMaker, AWS Bedrock, AWS Lambda, OpenSearch

Sep 2022 – **Machine Learning Engineer**, *LexisNexis (via Devoteam)*, Paris, France

Jun 2023 **Context:** Mission within the LegalTech division, in a 3-person MLOps squad, to scale ML products used by thousands of legal professionals worldwide.

- Set up a full MLOps platform, reducing deployment cycles by 41% and improving release reliability
- Refactored Python codebases to increase maintainability and reduce tech debt
- Built CI/CD pipelines (GitLab, ZenML, MLflow, DVC) for automated ML workflows
- Managed AWS infrastructure with Terraform (including SageMaker) for scalable deployment

Stack: Python, AWS, Terraform, GitLab, ZenML, DVC, MLflow

Jul 2021 – Jul 2022 **Data Engineer**, *Generali (via CGI)*, Paris, France

Context: Mission within the Risk & Compliance department (team of 5), focused on strengthening fraud prevention and AML/CTF compliance.

- Implemented AML/CTF fraud detection system, enhancing compliance and reducing false negatives
- Built PySpark/Hadoop pipelines to improve data quality across millions of transactions
- Automated fraud detection workflows, reducing manual workload by 24%
- Used Neo4j (CypherQL) to detect fraud networks and suspicious patterns

Stack: Python, PySpark, Hadoop, Neo4j, CypherQL

Mar 2021 – **Data Scientist**, *BPI France, AG2R La Mondiale (via CGI)*, Paris, France

Jun 2021 **Context:** Joint project on document analysis (OCR + NLP), within a 4-person innovation lab, aiming to automate administrative workflows.

- Developed ML/DL models for OCR of printed and handwritten text, reducing manual processing time by 34%
- Improved prediction accuracy via hyperparameter tuning (+11.8% F1-score)
- Conducted text similarity analysis to enhance classification and deduplication of documents

Stack: Python, TensorFlow

Oct 2018 – **Data Scientist**, *DGFiP (via CGI)*, France

Jun 2019 **Context:** Mission within the Ministry of Finance data department, supporting modernization of citizen-facing digital services with NLP automation.




- Built NLP models for text classification, topic modeling, and chatbots
- Automated information extraction pipelines, improving processing efficiency
- Developed Python web apps (Flask) and deployed scraping pipelines




Stack: Python, Flask, NLP (NLTK, Word2Vec, LDA, SpaCy, Gensim)

Education

2018 **Engineering Degree**, *École Centrale de Lyon*, Lyon, France

Certifications

 AWS Cloud Practitioner Certificate
 AWS Data Engineer Associate Certificate
 AWS Solutions Architect Associate Certificate

 AWS Machine Learning Specialty Certificate
 Google TensorFlow Developer Certificate
 DCA Docker Certified Associate

Languages

French: Native (Mother tongue) • English: Expert (C2 level)