

PROFESSIONAL SUMMARY

Senior Data Scientist with 9+ years shipping production AI/ML — forecasting, anomaly detection, and LLM applications — across enterprise SaaS, fintech, retail, and CPG. Built on additional ~24-year foundation in mission-critical infrastructure.

TECHNICAL EXPERTISE

Languages:	Python, R, SQL, Shell
ML/AI:	Anomaly Detection, Forecasting, Time Series, NLP, LLMs, Deep Learning
ML Libraries:	scikit-learn, SparkML, XGBoost, statsmodels, Prophet, PyTorch, HF Transformers
GenAI Stack:	LangChain, LlamaIndex, LoRA fine-tuning, RAG patterns, Vector DBs (FAISS, Pinecone)
Big Data:	PySpark, Hadoop, Hive, Impala, ETL
MLOps/DevOps:	MLflow, Docker, Kubernetes, FastAPI, CI/CD, model monitoring
Cloud & Data:	AWS, Azure, GCP, Redshift, Oracle DB

PROFESSIONAL EXPERIENCE

Senior Data Scientist -> Senior Principal Data Scientist @ NetSuite AI | Oracle (Spain)

Senior Principal Data Scientist 09/2024 - 02/2026

- Shipped invoice payment-date forecasting model to production (XGBoost, MLflow) across 42,000+ enterprise accounts — +8% cash-flow prediction accuracy, materially less manual finance review.
- Built predictive duration models for task and project scheduling — cut planning estimation error ~13% inside the NetSuite PM module.
- Architected end-to-end anomaly detection for financial transactions — ~68% recall at <5% false-positive rate, directly reducing customer risk exposure.
- Selected as 1 of ~35 engineers company-wide for the Oracle AI safety initiative; evaluated 15+ AI systems pre-release.
- Delivered two keynote sessions at SuiteWorld 2024 (Las Vegas): "Streamline the Period Close: Harnessing AI and ML to Improve Efficiency" and "Scaling ML Pipelines for Financial Exception Management."

Senior Data Scientist 06/2021 - 08/2024

- Fine-tuned open-source LLMs (T5-Large with LoRA) for internal document Q&A and text-enhancement PoC — directly informed Oracle's investment decisions on LLM-based product features.
- Built bank-record reconciliation models to automate bank-to-general-ledger transaction matching for enterprise customers.
- Led cross-functional teams on high-impact AI initiatives; presented AI/ML solutions and roadmaps to C-level executives.
- Delivered technical session at SuiteWorld 2023: "Anomaly Detection Techniques with Business Use Cases".

Senior Data Engineer - Marketing and Growth Team | Glovo 09/2020 - 06/2021 (Spain)

- Built end-to-end PySpark + S3 + Redshift pipelines from raw ingestion to analytics-ready consumption for a new internal system.
- Replaced a legacy data provider via a new external-integration pipeline — saved >€150K/month (~€1.8M/year) in vendor costs.
- Built an asynchronous cohort-reporting system integrating external sources with the data lake — enabled sharper marketing targeting and measurable churn reduction.
- Implemented data-quality monitoring across three solution endpoints to detect sync failures, improving reliability of business-critical reports.

EARLY DATA SCIENCE ROLES

JPMorgan Chase | Kimberly-Clark | Datastar | Freelance 09/2017 - 09/2019 (Argentina)

- JPMorgan Chase (Sr Data Scientist) — Built a predictive ML tool (linear regression, SVM, decision trees, random forests) to forecast additional resource needs during scheduled maintenance windows — targeted ~5% improvement in resolution times.
- Kimberly-Clark (Sr Data Scientist) — Engineered a predictive-maintenance model (LSTM, ARIMA) on Azure Databricks forecasting chemical spills in U.S. production lines (~\$1M/year savings per line); built an RFM + K-Means customer-segmentation product for Global Marketing (~2% annual revenue uplift); prototyped an NLP CV-ranking tool on Azure for Global HR.
- Datastar (DS & Analytics Specialist) — Engineered high-volume Sqoop ETL migrating 1.2B rows from Oracle to Hive in under 4 hours; built supervised/unsupervised R models for pharmaceutical clients.
- Freelance (DS Consultant) — Built predictive churn models in SparkML on Hadoop for a major telecom provider; architected data models and ETL workflows for the data-warehouse-to-production-model transition.

TECHNICAL LEADERSHIP & INFRASTRUCTURE FOUNDATION

Senior Principal Software Engineer | Oracle 2015 - 2017

Led LATAM performance engineering; delivered 200%+ resource-utilization gains and processor-architecture workshops.

Systems Principal Sales Consultant | Oracle 2013 - 2014

Architected SPARC / Engineered Systems solutions for C-level customers; contributed to \$15M revenue target.

Managing Partner | IT evoluxion 2006 - 2013

Built a Sun/Oracle/Red Hat services firm; team of 11 engineers, 25+ enterprise clients.

Early Career (Infrastructure & High Availability) | Sun Microsystems, Banco Santander

Mission-critical Unix (Solaris/AIX/HP-UX), HA clusters, DR and storage for major finance and utility customers.

EDUCATION

- **Master of Science (Statistics and Operations Research, Data Science specialization)** | Universitat Politècnica de Catalunya / 2019 - 2023
- **Bachelor of Science (Computer Science)** | Universidad de Belgrano / 1991 - 1996

PUBLICATIONS & PATENTS

Patents:

- Point-in-Time Relative Outlier Detection (Published)
- Generating Enhanced Queries Using Machine Learning Models (Issued)
- Generating Vector-Based Recommendations (Filed)
- Cross-Cluster Transaction Risk Assessment (Published)

Journal:

- Machine learning for anxiety and depression profiling and risk assessment in the aftermath of an emergency (Artificial Intelligence in Medicine - Vol 157, Nov 2024)