

Lewis NDAMBIRI

PROJECT EVIDENCE SNAPSHOT - Industrial AI, Manufacturing Analytics & Operations Systems

Portfolio | GitHub | LinkedIn | Nantes, France / Trento, Italy

POSITIONING - I am most excited by the space where AI, manufacturing, and operations meet: real constraints, messy data, people who need reliable answers, and systems that can be useful beyond a demo.

SELECTED EVIDENCE METRICS

100% CNC RAG page hit@5	839 manual chunks indexed	9.23% forecast MAPE	0.914 forecast R ²	31.77x HPC speedup	95% Fedrigoni output accuracy
-----------------------------------	-------------------------------------	-------------------------------	---	------------------------------	---

FLAGSHIP SYSTEMS

Offline CNC RAG Assistant Local, source-grounded AI for CNC operation and maintenance knowledge React/TypeScript + FastAPI + MQTT + InfluxDB + PostgreSQL + Docker	Built a privacy-preserving retrieval assistant using Haas CNC Mill NGC documentation as a realistic manufacturing corpus. The system indexes 839 manual chunks locally, combines vector, keyword, and procedural retrieval signals, and produces cited answers through an offline Ollama/Mistral setup. Outcome: Demonstrates applied LLM/RAG ability under industrial constraints: local deployment, traceability, operator safety, and documentation reliability.
FactoryPulse Remote real-time industrial production monitoring and control XGBoost, Pandas, Scikit-learn, Streamlit, Docker	Developed a containerized industrial monitoring platform with live MQTT telemetry, OEE and downtime analytics, alarms, audit logs, reports, role-based access, and secure remote command boundaries. Outcome: shows capability to bridge shop-floor realities with full-stack data products, operational dashboards, and governance-minded system design.
Demand Forecasting XGBoost, Pandas, Scikit-learn, Streamlit, Docker	End-to-end forecasting workflow on Rossmann-style retail data: 1,115 stores, 1M+ rows, time-aware holdout validation, shifted lag/rolling features, recursive future forecasting, and deployment-ready Streamlit/Docker artifacts. Outcome: 9.23% MAPE, 0.914 R², and 55% RMSE reduction versus naive baselines. Evidence of business-facing ML with validation discipline.
Industrial AI Challenge - Fedrigoni Python, Pandas, NumPy, XGBoost, CRISP-DM	Analyzed and modeled the paper-slitting process at Fedrigoni's Arco plant to support predictive production planning. The work focused on feature engineering, output forecasting, low-output anticipation, and workload-balancing logic under real manufacturing constraints. Outcome: practical shift from reactive production follow-up toward predictive planning support.
Maintenance Mode Prediction Python, Scikit-learn, Pandas, Statistics	Built an interpretable classifier for industrial machine operational states, then treated the project as a model-risk case: label audit, suspicious feature review, leakage investigation, and future-data availability checks. Outcome: stronger than a high accuracy score alone; shows governance mindset, reliability concerns, and industrial ML judgement.
Parallel Numerical Integration / HPC C, MPI, OpenMP, PBS, Linux HPC Cluster	Parallelized Romberg numerical integration across MPI, OpenMP, and hybrid configurations, then benchmarked scaling behavior across cluster placements. Outcome: 31.77x speedup on 32 cores with 99.3% efficiency; evidence of systems thinking beyond notebooks.

Links: Portfolio | GitHub | LinkedIn