Executive Summary — Applied ML

Sierra Napier

May 2026

Applied Machine Learning — Executive Summary

Sierra Napier | Data Scientist | AI Architect

What I Built

Three production ML systems using real public datasets — not mockups, not tutorials.

Project	Business Problem	Data Source	Key Result
Predictive Maintenance	Unplanned engine failures cost millions	NASA C-MAPSS: 200 engines, 21 sensors, 33K+ cycles	RUL prediction pipeline with degradation curves
NLP Classification	Manual document routing is slow	sklearn 20 Newsgroups: 18,846 real Usenet posts	67.87% accuracy (Naive Bayes) with TF-IDF + BERT comparison
Demand Forecasting	Bike-share inventory mismatches demand	UCI ML Repo: 17,379 hourly records	ARIMA + XGBoost ensemble for hourly prediction

---

Why This Matters

For Operations Teams: The predictive maintenance system demonstrates how to turn sensor data into actionable failure warnings. Instead of reactive repairs, you get 25+ cycles of advance notice.

For Document Processing: The NLP pipeline shows how to classify unstructured text at scale. Applied to FOIA requests, this same architecture achieved 100% topic routing accuracy.

For Inventory Management: The forecasting model proves that public datasets (bike sharing) can validate approaches that scale to private operational data (fleet management, retail inventory, staffing levels).

Tech Stack

Python · Pandas · Scikit-learn · XGBoost · PyTorch · Transformers · Plotly · Lifelines

---

Verification

All notebooks executed with real data. All figures generated from live datasets. No synthetic data.

Live Code: github.com/gosidehustlesisi/sierra-applied-ml
Interactive Portfolio: gosidehustlesisi.github.io/sierra-applied-ml/

"Theory without practice is empty. These projects are production-ready pipelines with real data, not tutorial notebooks."