

IEEE PES/IAS Power Africa 2026 Tutorial

AI for Energy Demand Forecasting and Predictive Maintenance

Solar PV Off-Grid System Design: Foundations and AI
Applications

Dr. Oluleke Babayomi



■ Tutorial Roadmap

- Introduction and machine-learning foundations
- Part A: AI for energy demand forecasting
- Part B: AI for predictive maintenance
- Hands-on Google Colab demonstration

01

Introduction and ML Foundations



■ **The Electrification Paradigm**

- AI data centres, digital infrastructure, and EVs are reshaping demand
- EV charging is driven by human behaviour, weather, and traffic
- Load is far more volatile than traditional industrial consumption
- Meeting spikes safely requires proactive planning, not reactive monitoring

■ From Grid to Off-Grid

- An off-grid mini-grid has no upstream network to absorb spikes
- Finite storage and diesel backup must be positioned ahead of demand
- EV charging is used as a proxy for a deferrable off-grid industrial load
- Every method carries over to a locally metered African site

■ **Time-Series Forecasting: Bridging the Gap**

- Analyse historical patterns to predict future sequential events
- The mathematical bridge between current capacity and future load
- Uncovers daily commuting cycles and weekly seasonality
- Drives dispatch, storage, and generator decisions

■ **What Machine Learning Is**

- Learns patterns from data, not from hand-written rules
- Supervised learning maps inputs (features) to known targets
- Two task types here: regression and anomaly detection
- No prior machine-learning background is assumed

■ **The Supervised Workflow and the Leakage Rule**

- Assemble, engineer features, split, fit, evaluate, deploy
- Split chronologically, never shuffle a time series
- Use only information known at the forecast origin
- Causality prevents data leakage: the recurring theme of both parts

■ **Model Families and Evaluation**

- Statistical: ARIMA captures autocorrelation and seasonality
- Tree ensembles: Random Forest and XGBoost, strong low-cost defaults
- Deep learning: convolutional and recurrent networks
- Metrics: MAE and RMSE; always beat a simple baseline first

■ **ML Framing for Predictive Maintenance**

- Gradual degradation: regression of a performance metric over time
- Abrupt or subtle faults: anomaly detection, baselines to learned methods
- Scarce fault labels: uncertainty-aware and augmentation methods
- Applied to PV telemetry in Part B

02

Part A: Energy Demand Forecasting



■ **Why Forecast at an Off-Grid Site**

- Finite battery storage; diesel generator as backup
- Pre-charge storage before peaks; schedule the generator early
- Fuel burn and battery cycling dominate operating cost
- Forecasting is decision support, not a dashboard

■ Load Context and Dataset

- UrbanEV, Shenzhen: one charging zone, 4344 hourly records
- September 2022 to February 2023; volume, price, weather
- Target: hourly dispensed energy volume (kWh)
- Strong diurnal and weekly seasonality in the series

■ Problem Formulation

- Predict the value h hours ahead from information at the origin
- Horizons reported: 1 h, 12 h, and 24 h
- Fair-comparison rule: the newest usable lag is lag h
- Any exogenous input must be known for the target hour


■ Feature Engineering

- Lags starting at the horizon: h , $h+1$, $h+2$, and a 24 h seasonal lag
- Calendar encodings: hour-of-day and day-of-week
- Newest lag at h means no value inside the window leaks in
- Weather excluded: future weather is unavailable at forecast time


■ **Model Taxonomy**

- Baseline: Seasonal-Naive, the value 24 h earlier
- Statistical: ARIMA
- Tree ensembles: Random Forest and XGBoost
- Deep learning: one-dimensional convolutional network (1D-CNN)

■ Results: MAE



■ Results: RMSE



■ Model Comparison: MAE by Horizon (kWh)

Model	1 h	12 h	24 h
Seasonal-Naive	12.01	12.01	12.01
ARIMA	2.94	11.05	11.38
Random Forest	3.20	12.15	11.96
XGBoost	3.11	11.85	11.63
1D-CNN	4.96	12.65	13.24

RMSE follows the same trend. Horizon choice dominates model choice.

■ Forecast Horizon: Error Grows

- Error grows with the horizon as recent-lag signal decays
- These are simple, short-term, illustrative forecasters
- Microgrid planning needs forecasts up to six months ahead
- Storage sizing, procurement, maintenance, capacity expansion

3 → 12

MAE (kWh),
1 h → 24 h

■ **Extending the Horizon**

- Exogenous and calendar drivers with their own forecasts
- Multi-seasonal decomposition (STL, TBATS, Prophet-style)
- Long-horizon deep models (N-BEATS, N-HiTS, TFT, DeepAR)
- Global and transfer-learning models across many sites
- Probabilistic forecasting (quantiles and intervals)
- Scenario and ensemble methods for planning

■ Evaluation and ROI

- Report MAE and RMSE in kWh; MAPE omitted (unstable near zero)
- Quantile outputs give prediction intervals
- Coverage and interval width matter most at long horizons
- Model choice is a cost decision: less fuel and battery cycling

03

Part B: Predictive Maintenance



■ **Why Predictive Maintenance**

- PV arrays degrade slowly and can also fail abruptly
- Reactive repair loses generation; blanket schedules waste labor
- Target intervention to the assets that need it
- Use sensor telemetry to detect drift and anomalies early

■ Dataset

- NREL PVDAQ System 10; hourly records, 2006 to 2023
- Plane-of-array irradiance, DC and AC power, voltages, currents
- Three redundant module-temperature sensors
- Openly accessible; substitutes for a local African site

■ **Failure Modes and Sensor Signatures**

- Soiling, partial shading, string open-circuits
- Inverter clipping and faults; gradual module degradation
- Each maps to a sensor signature
- Signatures motivate the features used later

■ **Field-Data Cleaning**

- Convert the NREL -99999 error flag to proper nulls
- Remove physically impossible values (negative power, irradiance)
- One consolidated cleaning step; no sentinel reaches the models
- Subset to top-of-hour readings to suppress minute-level noise

■ **Exploratory Analysis and Degradation View**

- Plane-of-array irradiance and DC power are near-perfectly correlated
- Thermal loss above 800 W/m²; visible inverter clipping ceiling
- Monthly trends separate aging from seasonal sunlight
- Density shifts toward lower output with age

■ **Anomaly Detection: Statistical Baseline**

- Three redundant back-sheet temperature sensors
- Flag pairwise disagreement above 5 degC in daylight
- About 10,600 of 35,983 daylight records flagged
- No training and no labels; often the right first tool

■ **Anomaly Detection: Advanced Methods**

- Autoencoders learn a normal-operation manifold
- High reconstruction error flags anomalies
- Uncertainty-aware variants report confidence in each flag
- Confidence matters when a false alarm triggers a site visit

■ Degradation and State-of-Health

- Weather-corrected performance index regressed on time
- Robust nameplate proxy; no hardcoded slope
- Single 17-year series: no detectable trend
- Use a literature rate of 0.6 %/yr for planning

$$R^2 \approx$$

$$0.03$$

No detectable
degradation trend

■ **Fault Classification and Augmentation**

- Classify fault types with 1D-CNN or LSTM classifiers
- Augment scarce fault data with generative methods
- Outlines the full predictive-maintenance landscape
- Concept stage in this tutorial


■ ROI: Predictive-Maintenance Framing

- Benefit = degradation recovery + downtime avoidance
- Scaled to a mini-grid using the measured specific yield
- 100 kW: net positive; break-even near 84 kW
- 1 MW strongly positive; downtime dominates the value

+8%

ROI at 100 kW
(break-even ~84
kW)

ROI: Predictive-Maintenance Framing



■ **Deployment Roadmap and Architecture**

- 0 to 3 months: instrument the site, collect a clean baseline
- 3 to 6 months: train and validate models offline
- 6 to 12 months: deploy edge inference and close the loop
- Cloud training with lightweight edge inference

04

Hands-On Google Colab Demo



■ **Session 3 Workflow**

- Runs the Part A forecasting pipeline end to end
- Setup, data load, visualization, feature engineering
- Chronological split; lags start at the horizon h
- Notebook pre-executed; offline copy as fallback

■ **Models, Live**

- Baselines: Seasonal-Naive and ARIMA
- XGBoost and Random Forest, direct per horizon
- 1D-CNN, multi-output, scaler fit on training data only
- MAE and RMSE reported by horizon

■ **Guided Exercise and Probabilistic Extension**

- Change the look-back window or the horizon set
- Retrain and compare MAE across horizons
- Discuss context versus compute and latency trade-offs
- Sketch a quantile output layer for prediction intervals

■ Takeaways and Open Problems

Takeaways

- Features (lags + calendar) carry most of the signal
- Establish a baseline; horizon often beats model choice
- Forecasts drive decisions, not dashboards

Open problems

- Long-horizon forecasting (weeks to six months)
- Calibrated probabilistic forecasting on edge hardware
- Transfer to data-scarce off-grid sites

Thank You

Questions and discussion

Dr. Oluleke Babayomi

babayomi@ieee.org | obabayomi.github.io

