PROJECT ABSTRACT

Electricity Demand Forecasting using LSTM

Topic ID: K-1212 | Branch: AI & Machine Learning | Type: Software | Availability: Built-to-order

01. Title

Electricity Demand Forecasting using LSTM — short-term load forecasting with a stacked LSTM sequence-to-sequence model.

02. Introduction / Background

Short-term load forecasting underpins grid operation: generators are committed, reserves scheduled and prices set on the basis of predicted demand. Errors are expensive in both directions — over-forecasting idles capacity, under-forecasting strains the system.

Residential demand is strongly structured (twin daily peaks, weekday/weekend shapes, seasonal drift) and therefore learnable. Recurrent networks, and LSTMs in particular, suit this task because load has long memory: this evening's shape depends on patterns stretching back days. This project trains a stacked LSTM on the UCI household power consumption dataset (2,075,259 minute-level records, 2006–2010) to map a 168-hour window to the next 24 hourly demand values, validated on strictly chronological splits.

03. Problem Statement

The core problem is honest short-term forecasting: predict the next 24 hours of demand from recent history with quantified uncertainty. The methodological trap — random train/test shuffling that leaks the future into training — is endemic in student work and produces fantasy accuracy. This project solves both the forecasting task and the evaluation problem: a seq-to-seq LSTM trained and tested on chronological splits, with prediction intervals that communicate uncertainty instead of hiding it.

04. Objectives

- Build an hourly demand series from the UCI minute-level household power data with gap handling.
- Engineer calendar and lag features; construct 168-hour input windows with 24-hour targets.
- Train a stacked LSTM (64 → 32 units, dropout 0.2, dense 24-output head) with Adam and early stopping.
- Validate chronologically (train ≤2009, test 2010) and report MAPE/RMSE/MAE per horizon (design target: 3.8% MAPE at 24 h).
- Render prediction intervals (±1 std) around forecasts in an interactive web demo with horizon control.
- Document LSTM theory, leakage avoidance and error analysis for the viva.

05. Proposed Methodology

- Resample the 2M+ minute records to hourly demand; clean gaps and outliers.
- Engineer hour/weekday/month and lag features; build (168 h × 6 features) → (24 h) sequence pairs.
- Chronological split: training up to end-2009, validation/test on 2010 — no shuffling, no leakage.
- Train the stacked LSTM with MSE loss, Adam optimizer, early stopping on validation MAPE.
- Evaluate held-out weeks once: MAPE, RMSE, MAE per horizon; error sliced by hour-of-day and season.
- Serve the trained weights in the web demo: latest window in, actual-vs-forecast with intervals out.

06. System Architecture / Working

Data pipeline → feature store → sequence builder → stacked LSTM → dense 24-head → evaluation harness → web demo. The two LSTM layers (64 units returning sequences, then 32) carry weekly context across the 168-hour window; dropout 0.2 regularizes; the dense head emits all 24 hourly steps at once so the forecast keeps a coherent daily shape. Exogenous calendar features are concatenated at each step.

Dataset: UCI 'Individual household electric power consumption' (2,075,259 records, Dec 2006–Nov 2010, Sceaux, France). Task: multivariate sequence regression. Model: stacked LSTM, ≈45k parameters. Metrics: MAPE/RMSE/MAE — design targets, never claimed as measured.

07. Hardware / Software Requirements

- Hardware: standard laptop/desktop, 8 GB RAM; CPU training sufficient (≈20–40 min); no GPU required.
- Software: Python 3.10+, TensorFlow/Keras, pandas, NumPy, Matplotlib, scikit-learn, Jupyter.
- Demo: any modern browser; single-file web app runs offline after download.

08. Expected Outcomes

- Trained stacked-LSTM forecaster with inference code and saved weights.
- Training notebook with loss curves, early-stopping diagnostics and per-horizon error tables (3.8% MAPE design target at 24 h).
- Interactive forecast demo: actual vs forecast, prediction intervals, 24/48 h horizon switch.
- Report, PPT and viva Q&A; on LSTM gates, seq-to-seq design, leakage and MAPE vs RMSE.

09. Applications

- Utility day-ahead planning: generator commitment and reserve scheduling.
- Microgrid and rooftop-solar dispatch: matching local generation to expected load.
- Demand-response programs: targeting flexibility offers at predicted peaks.

- Campus/commercial building energy management and load shifting.
- EV charging planning: anticipating evening charging peaks.

10. Future Scope

- Weather inputs (temperature, humidity) as exogenous regressors.
- Probabilistic forecasting: quantile heads instead of point + interval.
- Transformer / temporal-fusion architectures as a comparison study.
- Transfer learning: fine-tune the household model to new feeders or regions.
- Multi-site hierarchical forecasting with reconciliation.