



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Name: Mehul Tank

Student Id: 2410391

University: University of Wales Trinity Saint David Birmingham

Course: MSc Data and Artificial Intelligence

Module: Master's Project

Supervisor: Dr Abdallah Naser

Short Term Electricity Demand Prediction in Great Britain Using Machine Learning

Abstract

The aim of this project is to explore a machine learning framework for short term electricity demand forecasting in Great Britain, considering new challenges relating to increasingly embedded renewables, the increasing electrification of transport, and changing demand patterns. This study will develop temporal, meteorological, and operational features using three public, open datasets: Elexon BMRS FUELHH (generation by fuel type, half hourly), the ESPENI dataset (Zenodo), and Open Meteo hourly weather data. Forecasting models based on Long Short Term Memory (LSTM) networks, Extreme Gradient Boosting (XGBoost) models, and Meta's Prophet model will be developed in addition to the model being tested against one another using timeseries cross validation and evaluated using RMSE, MAE, and MAPE metrics. SHAP values for XGBoost and attention mechanisms for LSTM will be used to investigate model explainability.

It is expected that LSTM would outperform the other two models to capture complex nonlinear temporal dependencies in the demand, while XGBoost will produce the lowest mean error applied to more structured features. Finally, the Prophet model will provide a more interpretable seasonal trend that is advantageous for operational planning. Ultimately, the research will provide a comparative study of these advanced models, shedding light on the influences of weather and calendar variables on electricity consumption, while also presenting a reproducible, data driven forecasting pipeline for both grid operations and policy development in Great Britain.

Overview: Forecasting electricity demand is essential for achieving a proper balance between supply and demand in modern power systems. Accurate forecasts of short term demand between one hour and one day ahead serve the important purpose of allowing grid operators to dispatch generation, retain adequate levels of stability, and manage operating costs. In Great Britain, increasing renewable generation, the electrification of transport, and the evolving consumption patterns are leading to a more volatile electricity demand, which is putting pressure on traditional forecasting methods based on time series models.

Purpose: This research seeks to develop and assess a machine learning framework for short term electricity demand forecasting in Great Britain. This study combines

advanced models (Long Short Term Memory networks, Extreme Gradient Boosting, and a Prophet model from Meta's) and compares their respective predictive accuracy against real time data for assessment.

Method: This work will measure time series data from three open datasets: 1. Elexon BMRS FUELHH (generation by fuel type, half hourly) 2. ESPENI dataset (Zenodo) 3. Open Meteo hourly weather data. Feature selection will comprise temporal, meteorological, and operational features. The models will be trained using time series cross validation for training and will be evaluated against the real datasets using RMSE, MAE, and MAPE measures. Model explainability for XGBoost will be investigated with SHAP values and LSTM will be examined using attention.

Expected Results: LSTM is expected to provide greater performance than Prophet and standard statistical fixes for extreme nonlinear temporal dependencies, while XGBoost will provide the least mean error for tabular structured features. Prophet is expected to provide interpretable seasonal trends that are useful for operations based planning.

Overall Conclusion: The study will provide a comparative study of three state of the art algorithms that are post in depth GB demand forecasting, provide insights into the influence of weather and calendar effects, and provide a reproducible, data driven pipeline for both national grid operations and policy application.

1. Introduction

Forecasting electricity demand is a key aspect of operating and planning modern power systems. Short term forecasts (30 minutes to 24 hours ahead) are critical for system operators, such as National Grid ESO, to efficiently dispatch generation, reduce spinning reserve requirements, and support system reliability in the most economy [1], [2]. In Great Britain (GB), the climate transition is progressing with increased penetration of intermittent renewable generation sources such as wind and solar, introducing greater volatility and uncertainty in the balancing of the system [3]. In addition, the demand profile is changing due to novel drivers, including the charging of electric vehicles, evolving time of use tariffs, and shifting consumer behaviours driven by climate consciousness and energy efficiency principles [4].

Statistical models, such as Autoregressive Integrated Moving Average (ARIMA) and exponential smoothing, have long been regarded as suitable methods for forecasting electricity demand. However, the adoption of traditional statistical models for forecasting electricity demand has been problematic, since they are not easily able to capture the nonlinear, non stationary, multi factor interplay demand drivers seen in modern power systems, especially with a high renewable mix [5]. Machine learning (ML)

and deep learning (DL) methods present a new approach and techniques to learn temporal dependencies and relationships between exogenous variables [6], [7].

Recent research has shown that Long Short Term Memory (LSTM) networks are well suited for modelling long term temporal dependencies in load sequences, achieving superior performance compared to conventional time series approaches in many cases [7]. In the meantime, gradient boosting models (for example, XGBoost), offer a flexible and interpretable framework for dealing with multiple and heterogeneous input features [6], [8]. On the other hand, Prophet, an open source time series forecasting tool developed by Meta, adopts a decomposition based model, which additively models trend, seasonality, and holiday effects in a manner which is interpretable and very much used operationally [9], [10], [14].

Although short term load forecasting has been examined in detail globally, there has been little research effort that has investigated this specifically in Great Britain (GB) and used openly accessible and reproducible datasets. Much of the previous research has used proprietary utility datasets, therefore reducing transparency and the possibility for benchmarking is limited. This proposed research will address this gap in the literature by synthesizing three separate open datasets (Elexon BMRS FUELHH, ESPENI and Open Meteo) into a single modelling framework, to evaluate and compare the predictive performance of three leading models (LSTM, XGBoost and Prophet) for short term load prediction across Great Britain. By evaluating their performance in a systematic and unified manner an aim of this research will be to evaluate their ability for predicting short term load across Great Britain, improving reliability, transparency, and reproducibility in forecasting the British grid with more confidence, bolstering the transition of GB towards a resilient and decarbonized power system [3], [4], [28]–[30].

2. Research Background

Demand for electricity in Great Britain follows unique daily and weekly patterns influenced by human behaviour, weather, and socio economic activities. The balancing mechanism administered by National Grid ESO needs to be able to accurately forecast demand in a half hourly format to keep generation and consumption in balance in real time. In the UK, we have typically relied on regression and classical time series techniques, including structural and autoregressive models, to support modelling demand by including weather and calendar variables [11], [16]. However, as the UK accelerates decarbonisation and seeks to integrate more renewable generation, the variability produced by this transition means that we can see the emergence of non linearities and uncertainty in the forecasting models, which traditional methods are less adaptable to capture (see [4], [17] for more details).

In recent years, the field of energy forecasting has shifted to data driven and machine learning (ML) based techniques. For example, neural networks can learn complex, nonlinear mappings from input variables—temperature, wind speed, time of day—to electricity demand without making any explicit assumptions about the functional relationships underlying the mappings [12]. Gradient boosting models, including XGBoost, are gaining traction because they are interpretable, robust, and scalable to larger datasets [13]. Similarly, Prophet (developed by Meta) can automatically decompose time series into trend, seasonality, and holiday effects to produce accurate, interpretable forecasts with a minimal amount of parameter tuning [14].

These different types of forecasting are also closely aligned to the UK's wider ambition to digitalise its energy system, and to support data led decision making across system operations. See the National Grid ESO's Digitalisation Strategy which includes creating a "digital twin" of the national grid, which is essentially integrated data analytics environment to simulate and optimise behaviour of the system in close to real time [15].

This research will leverage recent advances in academia and industry to create a reproducible ML based forecasting pipeline leveraging open and transparent data sources. In using explainable and ethically responsible practices that adhere to the FAIR data principles (Findable, Accessible, Interoperable, and Reusable), the robust framework described below aims to emphasize methodological transparency, open science, and replicability in short term electricity demand forecasting [25].

3. Literature Review

The need to predict short term electric power utilization in Great Britain has become ever more important due to increased decarbonisation, simultaneous developments in renewable generation, and shifting patterns of consumption [3], [4]. Traditional, statistical models such as ARIMA and Exponential Smoothing methods have provided a foundation for understanding short term demand variations, but these models have limitations in explaining nonlinear and volatile patterns of demand variation [8], [16].

In recent times, machine learning (ML) methods of prediction, such as XGBoost and Random Forests, have improved upon existing prediction performance, especially in explaining features and in recording nonlinear patterns of demand [6], [19]. Recently, deep learning (DL) techniques such as Long Short Term Memory (LSTM) networks improved temporal forecasting performance in power demands by accurately capturing longer term dependencies in the data (in particular, in periodicity in demand) [7, 9]. Emerging datasets of open data, for example, datasets publicly available via Elexon

BMRS, ESPENI, and Open Meteo, will allow for forecasts coded to have reproducibility and explainability, and ultimately are aligned with the principles of FAIR Data [25, [28] [30].

3.1 Classical Approaches

In the area of short term load forecasting (STLF), statistical methods have always been fundamental, particularly using models such as Autoregressive Integrated Moving Average (ARIMA) and Exponential Smoothing techniques [16]. Taylor and McSharry [8] showcased how effective these classical methods were, by determining reasonably accurate predictions of the forecasting horizons of Great Britain's system demand. However, such models lack robustness toward structural change, abrupt regime change, and changing behaviour in demand time series. Kalman filtered state space models and multiple linear regression (MLR) have also been applied to demand forecasting in electric markets [11]; however, linearity, normality, and homoscedasticity assumptions limit their use and extrapolation.

3.2 Machine Learning Methods

A transformation in methodology has occurred in STLF with the emergence of machine learning (ML) methods. Support Vector Regression (SVR) [18], established in statistical learning theory, alongside ensemble methods such as random forests [19], yield improved performance through more powerful nonlinear modelling and feature interplay. Additionally, new gradient boosting algorithms, such as XGBoost, feature an ability to handle missing data (and missingness), allow for weak learners, and provide feature importance metrics to aid the field of interpretability [6], [13].

Hong et al. [20] and Chen and Guestrin [6] showed that with exogenous variables such as temperature and holidays, gradient boosted trees decreased the forecasting error rate by 30% over ARIMA models. These advantages have led to the widespread usage of XGBoost within the industry, government, and consulting spaces and have persisted because of its open source availability for strong scalability and transparency.

3.3 Deep Learning Methods

The emergence of deep learning (DL) has transformed time series forecasting with architectures that learn sequential dependencies directly from the data. Recurrent Neural Networks (RNNs) have successfully been employed in various forms of temporal data [1]. The improved architecture, Long Short Term Memory (LSTM) networks, reduces the vanishing gradient problem through gating mechanisms for long term temporal dependencies [7]. Ryu et al. [9] found that LSTM based models improved forecasting of energy related data accuracy by more than thirty percent, relative to standard feedforward neural networks.

More recently, Transformer based architectures have arisen as equally powerful alternatives for modelling sequential data and uniquely capture long range dependencies using self attention mechanisms [21]. While these frameworks are promising, they require a significant number of computational resources and volume of historical data, which make them less relevant to a small scale operational forecasting job. This study proposed the use of LSTM networks given the balanced consideration of predictive power and computational efficiency.

3.4 Hybrid and Ensemble Approaches

Hybrid and ensemble forecasting methods aim to combine the distinct advantages of different model classes. In general, statistical models have proven effective at capturing seasonal and trend components, while ML or DL models are generally better suited for learning nonlinear residual patterns [5], [22]. For example, Zhao et al. [22] developed a hybrid ARIMA–LSTM framework which examined China's power grid, which outperformed the two models existing separately. A similar hybridisation approach could potentially apply to GB forecasting, which has exhibited the same temporal and meteorological complexities.

Prophet, which is an additive time series model developed by Meta, is an interpretable benchmark model that decomposes time series into trend, seasonality and holiday components [10], [14]. Because it is interpretable and robust, it should serve as a useful benchmark for LSTM and XGBoost models, which are examples of "black box" models [12].

3.5 Feature Engineering and Data Sources

Feature engineering is important for improving prediction accuracy. Weather variables, including temperature, humidity, wind speed, and solar irradiance, are often the most impactful exogenous predictors of electric demand [31]. Reanalysis datasets such as ERA5 or weather APIs like Open Meteo can provide high resolution meteorological data on temperature, wind, and cloud cover that can be used as input in ML pipelines [24], [29].

Previous studies have found that lagged demand values, rolling averages, and categorical variables that illustrate whether the observations occurred on weekdays, weekends, or public holidays also improve the predictive performance of models [26]. The ESPENI dataset [28] that contains energy system variables along with weather variables collected from multiple European nations serves as an additional reference and provides useful contextual and correlation features for creating more generalisable models.

3.6 Research Gap

Although there is significant international coverage on ML based load forecasting, literature focusing on Great Britain is sparse and less cohesive compared to international literature. Most of the GB affiliated studies in the literature are based on proprietary data or data from a single source; this hinders transparency and reproducibility. As well, few studies systematically compare interpretable models, like Prophet [14], tree based models (e.g., XGBoost [6]), and sequential deep models (e.g., LSTM [7]) in a common experimental paradigm.

This research seeks to fill these gaps in the literature by establishing an open, reproducible forecasting pipeline that incorporates Elexon BMRS FUELHH, ESPENI, and Open Meteo datasets [28]– [30]. The open research approach follows many of the principles of open science and is compliant with the FAIR data principles (Findable, Accessible, Interoperable, Reusable) [25]; this work will contribute to the eventual advancement of reproducible, explainable and operationally relevant forecasting methods for the GB electricity system.

4. Methodology

The approach taken in this study will establish a data driven, reproducible process for short term electricity demand forecasting in Great Britain, which will be systematic, transparent, and consistent. The outline of the research design will clarify how the study will be arranged systematically for methodological rigour and scientific validity. Specific processes will be outlined for data gathering, preprocessing, feature engineering, and model generation. Evaluation methodology on the models' performance will also be outlined. The methodological processes will also seek to ensure reproducibility, scientific relevance to the grid context, and comparisons of three models (LSTM, XGBoost, and Prophet).

4.1 Research Design

This research adopts a quantitative experimental research design exploring three types of machine learning model: Long Short Term memory, Extreme Gradient Boosting (XGBoost), and Prophet, and their application to short term electricity demand forecasting within Great Britain. The methodology will be sequenced in a workflow approach consisting of data gathering, data integration, data preprocessing, feature engineering, model generation, model validation, model evaluation, etc. Each process will be designed to ensure clarity, transparency, reproducibility, and robustness of the forecasting framework.

4.2 Data Collection and Sources

The initial step of the project is related to data acquisition, which will consist of three publicly available data sources. As a first source, the Elexon BMRS FUELHH dataset is used to receive half hourly, megawatt (MW) generation data by fuel type, which enables the calculation of total demand data to develop an understanding of how the mix of generation relates to demand behaviour. A second dataset, the ESPENI dataset published through Zenodo, is used to receive electricity load, generation, and weather data from several European countries. This dataset serves as a secondary dataset for cross referring and strengthen the assessment of weather demand relationships. A third data source is Open Meteo hourly weather data, which will be provided through its application programming interface (API). The weather data will also be made up of an array of meteorological data, including rainfall, temperature, wind speed, solar radiation, clouds cover, and they will also be updated hourly.

The data will be collected over three years to ensure that sufficient seasonal and climate variability is collected to build demand models to sufficiently generalize across a range of demand scenarios. All datasets will be accessed through the proper open access means and securely stored for future processing and analysis.

4.3 Data Integration and Preprocessing

Once the data is collected, all the collected data will be combined using the synchronized timestamp. Each piece of data will be converted to Coordinated Universal Time (UTC) and update the time for British Summer Time (BST) 'to make sure they are consistent'. The whole data will also be then aggregated to be hourly.

During the preprocessing stage, we will conduct quality, and consistency checks on the data by rationally cleaning the data. Where there is missing weather data, we will estimate values using linear interpolation and for the demand values, we will forward fill the missing values. The outliers will be flagged using z scores (to identify values greater than three) but kept as labelled anomalies to capture odd conditions caused, for example, by any restrictions due to COVID 19 lockdowns. Once that is established, we will normalize the features to increase the performance of the models. MinMax scaling will be used for deep learning models, for example, 'LSTM', while tree models such as XGBoost will use the raw value since the model is invariant to if the monotonicity of the transformation is not changed.

4.4 Feature Engineering

Feature engineering will be an important aspect for enhancing accuracy and interpretability of models. Temporal features will comprise the following hour of the day, day of the week, month, and binary indicators for weekend and holiday. To maintain cyclical structure, time features will be encoded as sine and cosine transformations.

Lag features such as previous hour, previous day, and previous week demand values will be created to capture the autocorrelation that exists in the time series. Weather features will include temperature lags and degree hour indices and temperature humidity interaction terms to capture the effects of weather variability on electricity consumption. Operational features, such as the percentage of renewable generation versus total fuel mix, will be incorporated to indicate demand changes that embed variability due to weather or changes in renewable supply.

Together these engineered features will improve model ability to identify and learn temporal and environmental dependencies that could be complicated, with respect to short term electricity consumption.

4.5 Model Development

The following models will be explored for this research:

LSTM: The Keras/TensorFlow framework will be utilized for the implementation of the LSTM model as a sequence to one architecture. The model will comprise two stacked LSTM layers with hidden units of size 64 and 32, respectively, followed by a dropout layer with a dropout rate of 0.2 to reduce overfitting, and a dense output layer featuring a linear activation function. The model will be trained using the Adam optimizer with a learning rate of 0.001, and the loss function will be Mean Squared Error (MSE). These properties will allow LSTM to learn long term temporal dependencies in electricity demand sequences, while also modeling more complex nonlinear relationships arising from weather and time.

XGBoost: The XGBoost model will be built as a gradient boosting framework that can use shrinkage and regularization to build additive regression trees that avoid overfitting. Some hyperparameters of interest will include maximum depth, learning rate (eta), and column sampling ratio, and they will be tuned with Bayesian optimization. XGBoost will be used to model nonlinear relationships with the predictors, as well as generate SHAP based feature importance values to build model explainability.

Prophet: The Prophet model, developed by Meta, will act as an interpretable baseline that decomposes the time series into its trends, seasonal effects, and holiday effects using an additive approach. The trend component will use piecewise linear growth, while the seasonal effects will be included at daily, weekly, and yearly granularity. Holidays in the United Kingdom, from the official Gov.UK dataset, will be included to measure the effect of these calendar events on electricity demand. Prophet will serve as a diagnostic and interpretable benchmark for assessing the performance of the more complex LSTM and XGBoost models.

4.6 Tools, Environment, and Reproducibility

The analysis will take place in Python, using libraries such as pandas, scikit learn, TensorFlow/Keras, XGBoost and Prophet in the Jupyter Notebook environment. For version control, the project will utilize DVC and GitHub which would facilitate reproducibility and transparency during the research process. Data visualization will also be utilized to communicate temporal trends, comparative results of model performance, and interpretability results, using either Matplotlib or Plotly.

The experimental work itself will be conducted on a local workstation, with the option of using GPU acceleration to train deep learning models. This will provide efficient computation and allocation of resources while maintaining complete control over the the execution of experiments.

In summary, the methodological framework assures that all aspects of the research process from data preparation to model evaluation will be transparent, replicable, and consistent with open science practices.

5. Expected Results and Discussion

This section will detail the anticipated outcomes stemming from the use of the proposed forecasting framework. The discussion will forecast how the chosen models will perform based on their theoretical properties and prior evidence from the relevant literature. The section will also describe the expected results in comparative performance, practicability, and broader implications for operational planning more broadly for energy policy.

5.1 Anticipated Model Performance

The three machine learning models will be expected to have different relative strengths depending on the forecast horizon. It is expected the LSTM model to have the lowest Root Mean Squared Error (RMSE) for the multi hour horizons because of the model's ability to capture long range temporal dependencies in the time series data. The XGBoost model will be expected to perform very well with the shorter forecasting horizons (one to six hours ahead) and provide interpretable insight into key features, including the temperature, calendar variables, and lagged demand. With SHAP analysis, it is expected that the XGBoost model will show these features as the highest importance predictors.

The Prophet model will be expected to deliver comparable accuracy in daily forecasts while being highly interpretable. The fact that it can break the demand down into trend, seasonal, and holiday components will highlight repeated patterns and longer term behavioural patterns that can be used in strategic and operational planning.

5.2 Comparative Insights and Evaluation

It will be expected that all three models will outperform standard statistical forecasting methods previously applied to Great Britain's electricity demand such as ARIMA and exponential smoothing. Each model will bring distinctive strengths: LSTM will improve accuracy in dynamic temporal forecasting; the XGBoost method will offer interpretability and robustness with structured features; and Prophet will present transparent and simple analysis for policy applications.

The models will be assessed on RMSE, MAE, and MAPE through time series cross validation, which will make sure that the evaluation will be statistically robust and generalizable. The anticipated comparisons will show trade offs between interpretability and predictive performance, informing future research and operational use in grid management.

5.3 Expected Practical Contributions

This study is expected to yield useful insights for the National Grid ESO and energy policymakers in general. The research will produce a reproducible, open access forecasting pipeline that illustrates how transparent data driven models can support short term operational decision making. The research will identify a set of predictors including temperature, hour of the day and prior day's demand that are the drivers of electricity demand in most instances.

The study will also quantify how to trade off model interpretability with accuracy for different modelling approaches to inform the future incorporation of data driven decision support systems for energy forecasting. Ultimately, the study is expected to facilitate Great Britain's transition to a more resilient, more intelligent, and decarbonized power system that has academic and operational relevance.

6. Ethical Considerations

All the datasets for this research are open access and non personal which means that there would be no exposure to individual data privacy risks. Ethical practice will be aligned with the tenets of data integrity, reproducibility, and transparency:


- **Data Ethics:** All datasets used will have publicly licensed datasets (Elexon BMRS, ESPENI, Open Meteo) consistent with the FAIR principles.
- **Research Conduct:** All data providers will be fully acknowledged with citations; also, the code associated with the research will be openly available through a MIT licence.

- **Bias Mitigation:** All evaluation of the models will be aware of seasonality, outliers, and structural change to mitigate against misleading generalisations.
- **Academic Integrity:** All writing will be free of plagiarism, and any references will be conducted in IEEE style, complying with an application for ethical approval if required by our university ethics committee.

7. Project Plan and Timeline (4 Months)

The project will span **Sixteen weeks**, structured as shown below.

GANTT Chart



Week	Task	Deliverables / Milestones
1–2	Literature Review & Dataset Inspection	Confirmed data availability; initial exploratory data analysis (EDA) and dataset understanding
3–4	Data Cleaning & Preprocessing	Cleaned and aligned datasets; preprocessing scripts and handling of missing values
5–6	Feature Engineering	Engineered time based and weather related features; feature selection summary
7–8	Prophet Model Development	Prophet baseline model with trend, seasonal, and holiday components
9–10	XGBoost Model Development	Trained XGBoost model; hyperparameter tuning and SHAP feature importance plots
11–12	LSTM Model Development	Trained LSTM model; training/validation performance curves
13–14	Model Evaluation & Comparison	Comparative performance table (RMSE/MAE); visual plots for forecast comparison
15	Interpretation & Discussion	Discussion of findings; feature influence and model interpretability analysis
16	Final Write up & Submission	Completed dissertation; appendices and GitHub repository submission

8. References

- [1] H. S. Hippert, C. Pedreira, and R. C. Souza, “Neural networks for short term load forecasting: A review and evaluation,” *IEEE Trans. Power Syst.*, vol. 16, no. 1, pp. 44–55, 2001. [Online]. Available: <https://doi.org/10.1109/59.910780>
- [2] T. Hong and S. Fan, “Probabilistic electric load forecasting: A tutorial review,” *Int. J. Forecast.*, vol. 32, no. 3, pp. 914–938, 2016. [Online]. Available: <https://doi.org/10.1016/j.ijforecast.2015.11.011>
- [3] National Grid ESO, “Future Energy Scenarios 2024.” [Online]. Available: <https://www.neso.energy/search/content?search=future%20energy%20scenarios%202024>

- [4] Ofgem, “Decarbonisation Action Plan,” 2020. [Online]. Available: https://www.ofgem.gov.uk/publications/ofgems-decarbonisation-action-plan?utm_source=chatgpt.com
- [5] G. P. Zhang, “Time series forecasting using a hybrid ARIMA and neural network model,” *Neurocomputing*, vol. 50, pp. 159–175, 2003. [Online]. Available: [https://doi.org/10.1016/S0925-2312\(01\)00702-0](https://doi.org/10.1016/S0925-2312(01)00702-0)
- [6] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD*, 2016. [Online]. Available: <https://doi.org/10.1145/2939672.2939785>
- [7] S. Hochreiter and J. Schmidhuber, “Long short term memory,” *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997. [Online]. Available: <https://doi.org/10.1162/neco.1997.9.8.1735>
- [8] J. W. Taylor and P. E. McSharry, “Short term load forecasting methods: An evaluation,” *Int. J. Forecast.*, vol. 22, no. 1, pp. 1–17, 2006. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0169207005000907?utm_source=chatgpt.com
- [9] Y. Ryu et al., “Deep neural network based demand forecasting for electricity consumption,” *Energies*, vol. 10, no. 3, p. 381, 2017. [Online]. Available: <https://doi.org/10.3390/en10030381>
- [10] S. J. Taylor and B. Letham, “Forecasting at scale,” *PeerJ Preprints*, 2017. [Online]. Available: <https://doi.org/10.7287/peerj.preprints.3190v2>
- [11] A. Harvey and S. Koopman, *Structural Time Series Models*. Oxford Univ. Press, 2000. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118445112.stat05808?utm_source=chatgpt.com
- [12] J. Lago et al., “Forecasting day ahead electricity prices with deep learning,” *IEEE Trans. Power Syst.*, vol. 35, no. 1, pp. 1–10, 2020. [Online]. Available: https://arxiv.org/abs/2008.08004?utm_source=chatgpt.com
- [13] Y. Li et al., “Short term load forecasting using improved XGBoost model,” *Energy*, vol. 219, p. 119592, 2021. [Online]. Available: <https://doi.org/10.1016/j.energy.2020.119592>
- [14] Meta (Open Source Team), “Prophet forecasting tool,” 2024. [Online]. Available: <https://facebook.github.io/prophet/>
- [15] National Grid ESO, “Digitalisation Strategy 2022.” [Online]. Available: <https://www.neso.energy/search-content?search=Digitalisation+Strategy+2022>

- [16] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control*, 5th ed. Wiley, 2015. [Online]. Available: https://www.wiley.com/en_gb/Time+Series+Analysis%3A+Forecasting+and+Control%2C+5th+Edition_p9781118674918
- [17] R. Weron, “Electricity price forecasting: A review,” *Int. J. Forecast.*, vol. 30, pp. 1030–1081, 2014. [Online]. Available: <https://doi.org/10.1016/j.ijforecast.2014.08.008>
- [18] V. N. Vapnik, *The Nature of Statistical Learning Theory*. Springer, 2000. [Online]. Available: https://link.springer.com/book/10.1007/9781475732641?utm_source=chatgpt.com
- [19] L. Breiman, “Random forests,” *Mach. Learn.*, vol. 45, pp. 5–32, 2001. [Online]. Available: <https://doi.org/10.1023/A:1010933404324>
- [20] T. Hong et al., “Global energy forecasting competition 2012,” *Int. J. Forecast.*, vol. 30, no. 2, pp. 357–363, 2014. [Online]. Available: <https://doi.org/10.1016/j.ijforecast.2013.07.001>
- [21] J. Lim et al., “Temporal fusion transformers for interpretable multihorizon forecasting,” *Int. J. Forecast.*, vol. 37, no. 4, pp. 1748–1764, 2021. [Online]. Available: <https://doi.org/10.1016/j.ijforecast.2021.03.012>
- [22] J. Zhao et al., “Hybrid ARIMA–LSTM load forecasting,” *Energy AI*, vol. 1, p. 100013, 2020. [Online]. Available: <https://doi.org/10.1016/j.egyai.2020.100013>
- [23] P. Brockwell and R. Davis, *Time Series: Theory and Methods*, 2nd ed. Springer, 2016. [Online]. Available: <https://iopscience.iop.org/article/10.1088/17426596/2803/1/012002>
- [24] Copernicus Climate Change Service (ECMWF), “ERA5 reanalysis dataset,” 2024. [Online]. Available: <https://cds.climate.copernicus.eu/>
- [25] FAIR Principles, “GO FAIR Initiative,” 2024. [Online]. Available: <https://www.go-fair.org/fair-principles/>
- [26] Gov.UK, “UK Bank Holidays Dataset,” 2025. [Online]. Available: <https://www.gov.uk/bank-holidays>
- [27] Sheffield Solar PV Live, “UK Solar Generation Estimates,” 2024. [Online]. Available: <https://www.solar.sheffield.ac.uk/pvlive/>
- [28] J. Lago, “ESPENI: Electricity system and weather data integrated dataset,” *Zenodo*, 2020. [Online]. Available: <https://zenodo.org/record/3973345>
- [29] Open Meteo, “Open Meteo weather API / historical data,” Open Meteo, 2025. [Online]. Available: <https://open-meteo.com/>

[30] Elexon BMRS, “Half hourly generation outturn by fuel type (FUELHH) — BMRS API Documentation,” Elexon, 2025. [Online]. Available: <https://bmrs.elexon.co.uk/api/documentation/endpoint/datasets/FUELHH>

[31] National Grid ESO, “Historic Demand Data Portal,” 2025. [Online]. Available: https://www.neso.energy/data_portal/historic_demand_data