

Highlights

Recurrent Neural Networks with Linear Structures for Electricity Price Forecasting.

Souhir Ben Amor, Florian Ziel

- A Parallel-Branch Recurrent Neural Network model with Skip Connections is introduced.
- We combine linear expert models with two branches of recurrent neural networks.
- Our methodology is tested on the largest European electricity market data (2018-2025).
- The hybrid network architecture provides superior electricity price forecasts.
- In terms of the RMSE, it outperforms state-of-the-art benchmarks by over 12%.

Recurrent Neural Networks with Linear Structures for Electricity Price Forecasting.

Souhir Ben Amor^{a,*}, Florian Ziel^a

*^aChair of Data Science in Energy and Environment
University of Duisburg-Essen
Germany*

Abstract

We present a novel recurrent neural network architecture specifically designed for day-ahead electricity price forecasting, aimed at improving short-term decision-making and operational management in energy systems. Our combined forecasting model embeds linear structures, such as expert models and Kalman filters, into recurrent networks, enabling efficient computation and enhanced interpretability. The design leverages the strengths of both linear and non-linear model structures, allowing it to capture all relevant stylized price characteristics in power markets, including calendar and autoregressive effects, as well as influences from load, renewable energy, and related fuel and carbon markets. For empirical testing, we use hourly data from the largest European electricity market spanning 2018 to 2025 in a comprehensive forecasting study, comparing our model against state-of-the-art approaches, particularly high-dimensional linear and neural network models. The proposed model achieves approximately 12% higher accuracy than leading benchmarks. We evaluate the contributions of the interpretable model components and conclude on the impact of combining linear and non-linear structures.

Keywords: Recurrent Neural Networks, Electricity Price Forecasting, linear expert model, Kalman Filter, Parallel Branches, Skip Connections.

1. Introduction, Motivation

Since the 1990s, the deregulation of the electricity market, as well as structural reforms, has resulted in radical changes in global markets. More precisely, a competitive market replaced the traditional vertically integrated electricity market. Therefore, like other commodities, electricity is today traded under competitive instructions using spot and derivative contracts ([Shahidehpour et al. \(2003\)](#)). For that purpose, companies that trade in the electricity market need robust price prediction methods to remain competitive.

Forecasting electricity prices is an essential and challenging issue for all market players due to its inherent features and the wide consumption of electricity in our modern society. Indeed, accurate electricity price forecasting is crucial for building risk management and formulating reasonable competition strategies. Specifically, producers and traders can develop bidding strategies to minimize risks, maximize profits,, and allocate purchases between spot prices and long-term bilateral contracts. In addition, electricity

*Corresponding author

Email addresses: souhir.benamor@uni-due.de (Souhir Ben Amor), florian.ziel@uni-due.de (Florian Ziel)

producers can rely on the forecast results to optimize unit production, while electricity consumers can adopt the forecast results to optimize their purchase portfolio (Abedinia et al. (2015)).

Therefore, price forecasting becomes an essential field of management, as future price change is a challenging issue for all market players. In fact, an accurate price forecasting represents the basis for the decision makers to propose an optimal decision and formulate a reasonable plan to reduce market risk and consequently maximize the social and economic management benefits (Hao et al. (2020); Zhao et al. (2014)). Moreover, considering that electric power is a clean and promising energy source, it plays an essential role in our daily lives since it is considered environmentally friendly compared to other traditional energy sources, and hence it is indispensable in the economic sector (Jiang et al. (2019); Wang et al. (2017)). However, regarding the specific characteristics of electricity and the uncertainty of market and bidding strategies, electricity price forecasts become more complex. Indeed, in contrast to other commodities, the impossibility of storing the electricity in order to use it in the future makes its prices more difficult to forecast. In other words, the electricity production and consumption should be taken simultaneously, which caused a high level of complexity and ambiguity in electricity markets.

This specific behaviour exhibits specific characteristics of the electricity price time series; such as high volatility, high frequency, non-constant mean and variance, a high percentage of unusual prices, unexpected price jumps or spikes quickly decay (associated with shock price elasticity demand and supply) (Nogales et al. (2002), Kuo & Huang (2018)).

Moreover, the seasonal behaviour of the prices is a direct result of demand fluctuations, which commonly arise from deterministic conditions (such as the weekly business hours and the number of daylight hours) or climate conditions (like precipitation levels and temperature).

Furthermore, there is an increasing incorporation of renewable energy sources and the development of smart grids which is making the Electricity Price Forecasting (EPF) more challenging than it already is. Electricity grids are kept stable thanks to spot markets, which include the Day-Ahead Market, Intra-Day Market, and Balancing Market. It is noteworthy to mention that often, the high price volatility created by renewable generation forecast errors presents significant challenges as these sources are often intermittent and subject to variability in production (Ortner & Totschnig (2019)).

Accurate forecasting is vital in strategically managing energy production, optimizing investment plans, and policy decisions aimed at risk reduction. In this regard, precise forecasting supports sustainable overall financial stability of the organization. In addition to achieving financial security, precise EPFs are instrumental in achieving climate goals by reducing emissions and aiding the incorporation of renewable energy into grids (O'Connor et al. (2025)). All these various reasons make the prediction of electricity prices very difficult. This in turn has significantly enhanced research efforts toward modeling and forecasting spot electricity prices.

Over the past few decades, various forecasting techniques have been developed. Statistical models (Muniain & Ziel (2020)), and Artificial Intelligence methods (Lago et al. (2021)) representing two major paradigms. Traditional statistical models, such as time series models, multivariate regression, and econometric approaches (Misiorek et al. (2006), Ziel & Weron (2018)), have been extensively applied to EPF due to their ability to capture linear relationships and underlying seasonality in electricity markets

(Ben Amor et al. (2018)). Misiorek et al. (2006) proposed an the linear expert model for spot price electricity forecasting which is considered as linear state-of-the-art model for its interpretability and robustness, particularly because they enhance forecasting efficiency by explicitly incorporating fundamental factors like load demand, fuel prices and renewable energy generation. These models are widely adopted later in various EPF researches (Uniejewski et al. (2016), and Ziel (2016)).

Although the electricity market is characterized by complex price formation processes and nonlinear interactions between fundamental factors, purely linear or parametric statistical models are often limited in effectiveness, particularly when market conditions are volatile or extreme. To address these limitations, Artificial Intelligence (AI) methods, particularly Machine Learning (ML) and Deep Learning (DL) techniques, have gained prominence in EPF. These methods excel at capturing nonlinear relationships and complex patterns in large datasets, making them suitable for modeling the intricate dynamics of electricity prices. Several deep learning models are considered for Day ahead EPF (Pesenti & O’Sullivan (2025), López et al. (2025), and Lago et al. (2018)), and they are considered as state of the art models used as benchmarks in the field (Lago et al. (2021)). In particular, RNNs have shown promise for modeling temporal dependencies and nonlinear dynamics present in electricity markets (Castello & Resta (2025), Yan et al. (2025), and Yang & Schell (2024)).

EPF has evolved in recent years, but there is still an ongoing debate about which approach is the best, since linear and nonlinear models both have distinct advantages. There have been studies that combine both approaches into a pure econometric hybrid model (Ben Amor et al. (2018), and Zhang et al. (2020)). In this context, our paper integrates an expert model (Misiorek et al. (2006), Uniejewski et al. (2016), and Ziel (2016)) into the architecture of a recurrent neural network (RNN) to take advantage of both models’ strengths in forecasting day-ahead electricity prices, we implement also a Kalman filter branch. There are many directions in which this work contributes to the existing literature, which can be summarized in the following way:

- We propose a novel parallel-branch RNN architecture with skip connections (linear expert model), which allows the model to capture both linear and nonlinear relationships in the data. To the best of our knowledge, this is the first time such an architecture with such model components is applied to day-ahead electricity price forecasting.
- On one hand we use the Kalman filter to replace the RNN in the architecture to highlight the importance of non-linearity in the RNN component by comparing it with a Kalman filter configuration results. On the other hand comparing the Kalman filter configuration with the linear expert model results to highlight the importance of temporal dynamics in the linear expert model.
- We demonstrate the effectiveness of our combined forecasting model on real-world data from the largest European electricity market between 2018 and 2025, which is characterised by the presence of different data regimes mainly due to the energy crisis in 2022.
- We implement our model using PyTorch. The code is available on GitHub ¹, making it accessible

¹<https://github.com/souhirbenamor/RNN-for-EPF>

for further research and development in the field of electricity price forecasting.

- We provide a deep analysis of the model hyperparameters selected during the optimization process, and explain how they impact forecasting accuracy.
- We provide a comprehensive evaluation of our model’s performance, including comparisons with existing state-of-the-art models, and discuss the implications of our findings for future research in electricity price forecasting.
- We also provide a forecast decomposition analysis to understand the contributions of different components of our model to the overall forecasting performance.

The rest of the paper is organized as follows: Section 2 provides a literature review of existing methods for electricity price forecasting. Section 3 describes the data and study design. Section 4 presents the methodology, including the architecture of our proposed models. Section 5 discusses the training procedure and hyperparameter optimization. Section 6 discusses the results and performance evaluation of our model. Finally, Section 7 concludes the paper and outlines future research directions.

2. Literature Review

Different forecasting techniques have been proposed over the past few decades to solve the aforementioned management challenge. EPF literature can be categorized into five main categories: (1) Fundamental methods (Kanamura & Bunn (2022), Paschalidou & Thomaidis (2025), and Tselika et al. (2024)) that model electricity price dynamics based on fundamental factors (e.g. demand, loads, and weather conditions). (2) Production cost models for long-term forecasting. (3) Reduced-form models, such as regime-switching models (Xiong & Mamon (2019), de Castro Matias & Tabak (2025)) which explain price dynamics through seasonal volatility, spikes, mean reversion, and correlations between commodity prices. In addition, (4) A substantial amount of literature focuses on quantitative or statistical analysis (Chen et al. (2025)) to develop modeling frameworks based on the values of historical prices and/or exogenous variables or regressors, besides probabilistic models (Marcjasz et al. (2023)). (5) Artificial intelligence-based (or non-parametric) methods (Ko et al. (2020)). (6) Hybrid approaches, which combine techniques from two or more groups listed above (Zhang et al. (2020), Zhang et al. (2019), and Amor et al. (2024)). Due to the fact that statistical and machine learning methods and their combined models have shown to yield the best results (Lago et al. (2021) and Weron (2014)), this review pays particular attention to them and the related benchmarking model will also be centered around them.

To describe electricity prices’ linear characteristics, statistical models use mathematical techniques combined with historical prices or price-related information to predict current prices. Among these models are the autoregressive moving average (ARMA) (Chu (2009)), autoregressive integrated moving average (ARIMA) (Ramos et al. (2015) and Zhao et al. (2017)), seasonal autoregressive integrated moving average (SARIMA) (Camara et al. (2016); Mohamed et al. (2010) and Soares & Medeiros (2008)), and the generalized autoregressive conditional heteroscedasticity (GARCH) (Garcia et al. (2005)). The ARIMA (or SARIMA) models have been coupled with GARCH models in some forecasting studies to

simultaneously model the conditional mean and the conditional variance (volatility) (Tan et al. (2010), Kumar et al. (2017)).

In addition to univariate baseline models, multivariate approaches incorporating fundamental market data are widely used in EPF, such as ARX and SARIMAX (X refers to regressors) (Mulla et al. (2024) and Misiorek et al. (2006)). In numerous studies, researchers have compared univariate and multivariate models. They show that the inclusion of exogenous variables enhances forecasting accuracy (Gianfreda et al. (2020), Ziel & Weron (2018)). Building on these findings, so-called "expert" models have emerged. In this model, domain knowledge and market fundamentals guide the selection of regressors. This model class is based on a parsimonious autoregressive structure, originally proposed by Misiorek et al. (2006). Since their introduction, expert-type models have become a central approach to electricity price forecasting (EPF). The literature widely documents their adoption (Uniejewski et al. (2016), Maciejowska & Weron (2015), and Weron & Misiorek (2008)). For example, Nowotarski & Weron (2018) benchmarks a classical ARX expert model and its multi-day extension (mARX) as primary forecasting tools for probabilistic day-ahead price prediction. They demonstrated competitive results for both point and interval forecasts using the GEFCom2014 dataset. The models incorporate expert-chosen lag prices, minimum price statistics, and load forecasts, sometimes outperforming neural networks and data-driven approaches. Ziel (2016) develops a LASSO regularisation expert-type autoregressive model that captures cross-hour and intraday variations in electricity prices. Applied to multiple European markets, the model achieves robust forecasting accuracy, outperforming traditional expert and benchmark models. An examination of 12 power markets was conducted in Ziel & Weron (2018) by comparing both univariate and multivariate model frameworks for estimating day-ahead electricity prices. Various expert-type and high-dimensional autoregressive models were assessed, and they found that performance varied across markets and contexts rather than being uniformly superior. In addition, the authors demonstrate that combining forecasts from both approaches can improve predictive accuracy and provide useful advice on selecting variables for lasso-type models. A recent study in Ghelasi & Ziel (2025) developed expert-type models for forecasting electricity prices in the German market over short, mid-, and long-term horizons. A comprehensive overview of the literature concludes that despite their state-of-the-art status for day-ahead forecasting, expert models based on autoregressive terms, calendar variables, and market fundamentals deteriorate over longer forecasting horizons. According to the authors, these challenges can be addressed by constraining model coefficients based on energy-economic theory and incorporating seasonal forecasts of key regressors. Their findings show that regularised linear models with carefully engineered constraints and inputs perform transparently and robustly for both operational and scenario-based long-term forecasting compared to machine learning approaches.

Given that these approaches are built on prior knowledge determined by experts, following Uniejewski et al. (2016) and Ziel & Weron (2018), we refer to them as expert models. Although these models can be interpreted, they often fail to capture the nonlinear dynamics in electricity markets (Chan et al. (2012)) as well as the high-dimensional interactions (Aggarwal et al. (2009) and Weron (2014)), which can lead to inaccurate electricity price forecasting.

Due to the limitations outlined above and to consider the features characterising electricity price

dynamics, ML models have gained increasing attention in recent decades. Most publications in the EPF field adopt these methods because they have shown excellent modelling performances due to their ability to handle unstable nonlinear dynamics in high-dimensional data, making them suitable for EPF (Hornik et al. (1989)).

As a result, a wide range of approaches has been developed that can be applied to complex dynamic systems and achieve accurate power price predictions (Castello & Resta (2025)). This includes Artificial neural network (ANN) (Lin et al. (2010), Panapakidis & Dagoumas (2016); Sandhu et al. (2016), Ortiz et al. (2016), and Keles et al. (2016)), fuzzy neural network (FNN) (Hung et al. (2014)), weighted nearest neighbors (WNN) (Lora et al. (2007)), and adaptive wavelet neural network (AWNN) (Pindoriya et al. (2008), Abedinia et al. (2015)), Feed-forward neural network (FFNN) (Anbazhagan & Kumarapan (2014)), support vector machine (SVM) (Yan & Chowdhury (2014)), and Random forest regressor (Tschora et al. (2022)).

As opposed to statistical approaches, ML does not make any assumptions about the functional form or statistical properties of the data set under consideration. Rather, nonlinear optimisation techniques are employed. Due to these features, they are more effective, accurate, and therefore popular for forecasting.

It has been demonstrated in recent EPF works that DL methods perform better than ML or statistical models (Lago et al. (2018)). This is attributed to the efficiency with which DL models map complex nonlinear relationships among price time series and external regressors (Li et al. (2022)).

To illustrate, Lago et al. (2021) introduced state-of-the-art deep neural network and statistical models evaluated on five European markets for day-ahead electricity price forecasting. They demonstrated that well-optimised DNNs often attain competitive or high accuracy. To promote comparable research on the subject, they provided both methodological guidance and a replicable benchmark framework. Recently, Yan et al. (2025) introduced a time-series decomposition-forecasting framework based on recurrent neural networks (RNN) for electricity prices. To break down time series data into trend, seasonal, and residual components, a hierarchical RNN-based time series decomposition model is presented. The time-series forecast is created by combining the predictions made by the RNN-based forecasting models for each component series. The suggested model is a promising and successful forecasting model, as shown by experimental findings.

The use of advanced RNN structures, such as LSTMs and gated recurrent units (GRUs) (Meng et al. (2022)) networks, has proven much more accurate when analysing complex nonlinear time sequences (Graves (2013) and Chung et al. (2014)), such as electricity prices. In Peng et al. (2018), an LSTM with a differential evolution algorithm has been used to predict electricity prices. To verify the performance of the proposed model, experiments are conducted under electricity price scenarios & Austria, New South Wales, and France. One example is Lago et al. (2018), which combines GRU and a DNN network as well as LSTM and DNN for day-ahead electricity price forecasting in Belgium. The DL algorithms were shown to be more accurate. In Lehna et al. (2022), a Convolutional Neural Network (CNN) combined with an LSTM model (resulting in a CNN-LSTM model) is used as a hybrid model for electricity price forecasting. To performance, they incorporated common external factors, including wind speed, average solar radiation, fuel and emission prices, and load. They suggested that the combined model leads to even

better predictions of electricity spot prices. It is important to emphasise that linear problems are difficult to handle (Tang et al. (2014)). Therefore, we cannot effectively extract all features present in pricing and achieve reliable forecasting results if we only model linearity. In this vein, theoretical and empirical research has shown that combining several models can be a useful strategy to increase forecast accuracy and mitigate the risk associated with the inadequacy of individual models. Integrating the advantages of each method to improve is the primary objective of combining different methods.

The existing hybrid models can be classified into three categories; For the first category, the original electricity price is decomposed into linear and non-linear components, and a single appropriate model predicts each component. Thereafter, the final predicted price is the sum of the two predicted components. In this category, the most adopted approach is to combine the ARIMA and the ANN models (Khashei & Bijari (2010); Tseng et al. (2002); Valenzuela et al. (2008); Valenzuela et al. (2008)). Despite that these hybrid models have achieved good forecasting results, they suppose that the nonlinear relations exist only in the residuals, and the linear and non-linear components can be modelled separately (Ben Amor et al. (2018)). Hence, the electricity price features cannot be well estimated.

To avoid this inefficiency, the second kind of hybrid model is proposed. Firstly, different models are applied separately to forecast the electricity prices; accordingly, different existing features can be captured by the suitable model. Secondly, an optimisation algorithm is adopted to compute each model's weight coefficient. The final forecasting results are the product of each model's prediction results and its weight coefficient (Zhang et al. (2022), Hussain et al. (2024), Ben Amor et al. (2024) and Lago et al. (2021)). To exemplify, Nowotarski et al. (2014) used twelve different individual models to predict electricity price separately, then identified weight coefficients of each model. Niu & Wang (2019) and Darudi et al. (2015) applied auto-regressive moving average, artificial neural network to predict electricity price; the weight coefficients of each model are identified. Although the accuracy of the second category of the hybrid model is enhanced compared to the first one, computing each model's weight coefficient is challenging, which can affect the forecasting accuracy.

To overcome the above-mentioned limitation, some researchers suggested a third kind of hybrid model. For this category, the electricity price is first decomposed into several components using a signal decomposition algorithm. Thereafter, an appropriate model is attributed to each component. Then, the final forecasting results are the sum of each component's forecasting results. Among the existing signal decomposition approaches, discrete wavelet transform, empirical mode decomposition, Variational mode decomposition are the popular methods for signal decomposition (Krishna Prakash et al. (2023), Zhang et al. (2020), Wang & Wang (2021), and Paraschiv et al. (2023)). To exemplify, Xiong & Qing (2023) proposed a new hybrid forecasting framework to improve the forecasting accuracy of day-ahead electricity prices. The proposed model consists of three strategies. First, an adaptive copula-based feature selection algorithm is proposed for selecting model input features. Second, a new method of signal decomposition technique for the EPF field is proposed based on a decomposition denoising strategy. Third, a Bayesian optimisation and hyperband optimised LSTM model is used to improve the effect of hyperparameter settings on the prediction results. The effectiveness of the different methods was validated using five datasets set up for the PJM electricity market, and the results indicated that the proposed hybrid algorithm is

more effective and practical for day-ahead EPF.

Despite the third kind of hybrid model having demonstrated good performance, the choice of the decomposition approach and the appropriate single model is not suitable for the electricity price, which makes it difficult to further improve the accuracy and stability of the hybrid model. Moreover, decomposition techniques are dependent on parameters, e.g. mode numbers, and are thereby vulnerable to configuration choices and less adaptable to changes in market conditions. This results in suboptimal decomposition, particularly when price spikes are expected (which is very common in electricity markets). By ignoring these factors, the decomposition fails to achieve its intended purpose, resulting in inaccurate as well as unstable forecasts (Luo & Shao (2024), and Chai et al. (2024)).

To develop a robust model for EPF, it is essential to first review the existing literature:

Previous studies revealed several gaps. First, expert-based statistical models are rarely compared directly with advanced ML approaches (Nowotarski & Weron (2018) and Ziel (2016)). In addition, to the best of our knowledge, no prior studies have explicitly integrated linear expert models and ML-based models together to exploit their complementary linear and nonlinear capabilities.

Second, most ML and DL research in EPF has concentrated on two main directions: (1) benchmarking the performance of various DL architectures (Mujeeb et al. (2019)), and (2) designing hybrid DL models that combine multiple architectures, such as convolutional neural networks (CNNs) and LSTM networks (Ahmad et al. (2019); Kuo & Huang (2018); Xie et al. (2018)). However, these studies typically omit comparisons with established state-of-the-art statistical models. For instance, (Ugurlu et al. (2018)) investigated RNNs for electricity price prediction, but restricted the analysis to a single market and benchmarked only against relatively simple statistical models such as seasonal ARIMA and Markov regime-switching. Although other DL methods were considered, the study did not evaluate their performance against simpler ML models.

While advanced RNN variants such as LSTMs are commonly used in EPF (Meng et al. (2022)), there is no conclusive evidence that these increasingly complex architectures consistently produce better forecasts. Particularly, comparisons with the basic Elman RNN are often absent, even though standard DNNs remain widely used and are considered state-of-the-art in EPF (Chinnathambi et al. (2018) and Lago et al. (2021)). Similar findings have emerged in financial forecasting, where Dautel et al. (2020) compared LSTM, GRU, simple RNN, and FNN for exchange rate forecasting and found that while deep models can perform well, they are also hard to train and tune. In some cases, simpler architectures matched or even outperformed more complex ones, especially in terms of trading profitability.

Motivated by these observations, we propose an EPF framework that combines an expert model with a standard Elman RNN and Kalman filter model. Using this design, we can assess each component's contribution as a stand-alone and exploit the strengths of both approaches.

Thirdly, NNs have one major disadvantage despite their advantages. It involves optimising a nonlinear, nonconvex loss function over a large parameter space, which can lead to local minima, slow convergence, or reduced robustness. With many weights and layers, overfitting becomes a concern since noise is incorporated into the learned patterns. Despite numerous mitigation methods, they remain challenging (Bejani & Ghatee (2021) and Baldassi et al. (2023)).

This is addressed in our approach through adaptive hyperparameter optimisation via Optuna, which allows the model to explore and select optimal configurations, learning rates, architectural settings, and regularisation parameters without manual trial and error. This, along with rolling-window evaluation, ensures that the model remains adaptable, retraining on the most recent market conditions and maintaining strong performance over time.

Finally, in the broader literature, many hybrid EPF models adopt highly complex designs, incorporating multiple algorithms for data decomposition, feature selection, clustering, and single forecast models whose forecasts are combined (Singh et al. (2018), Bento et al. (2018), and Xiao et al. (2017)). A further problem is that these studies rarely assess the impact of substituting different variants of hybrid components, leaving it unclear whether each module is relevant.

This is addressed in our method by integrating the expert model with a skip connection within the RNN architecture, with the final output being the sum of both components' forecasts. In spite of the simplicity of this structure, it retains the ability to capture both linear and nonlinear dynamics, making it easier to train and interpret.

Its worth noting that the Elman RNN is preferable to more complex LSTM-based models in short-term predictions of stochastic processes because the additional complexity of sophisticated architectures often yields marginal gains that cannot be statistically distinguishable Hewamalage et al. (2023).

The network is extended with an additional branch to leverage the time-dependence captured by the RNN under a linear activation function, resulting in a Kalman filter configuration. Our dataset spans 2018 to early 2025, covering both stable periods with low price volatility and highly volatile phases, including the COVID-19 pandemic, the global energy crisis of 2021–2023, and the post-crisis years. Through this dual-period coverage, our framework is tested under diverse and challenging market conditions.

3. Data and Study Design

The purpose of this paper is to forecast next-day electricity prices in the German electricity market using a multivariate deep learning model that predicts all 24-hourly prices simultaneously (as indicated in the network architecture in Figure 4). This setup mirrors the European day-ahead market, which determines hourly prices in a single auction each day. Specifically, in the day-ahead electricity market, market participants submit binding bids for each of the 24 hours of the following day ($T + 1$) by noon on day T . Afterwards, the market-clearing price (MCP) is determined at the point of intersection of the aggregated supply and demand curves, forming a uniform-priced auction. As the same information set is available for all 24 hours, we use a uniform feature set, incorporating multi-output neural architectures into our joint forecasting strategy.

The data is sourced from the ENTSO-E Transparency Platform (ENTSO-E)² and contains hourly observations of electricity prices, day-ahead forecasts of load, wind (on and offshore), and solar generation. Daily fuel prices (coal, gas, oil) and carbon emission allowances (EUA) are also included and they are

²<https://www.entsoe.eu/>

sourced from Datastream platform ³. This dataset covers the period between 10 October 2018 and 13 January 2025, and has been preprocessed to ensure time consistency and proper alignment.

Figure 1 plots the electricity price and the other above mentioned features used to predict it. According to the plot, the energy crisis led to record-high German power prices in 2022. This could be assumed to be derived from prices for fuels (particularly gas and coal), which are used to generate power Ghelasi & Ziel (2024). In this way, when gas and coal prices increase, electricity prices also rise.




Figure 1: Multi-panel time series for Germany’s day-ahead electricity price and related features (2018–2025). The panels show: Electricity Price (EUR/MWh), Day-Ahead Load Forecasts, Wind Onshore Forecasts, Wind Offshore Forecasts, PV (Solar) Forecasts, and fuel prices comprising Coal, Natural Gas, Brent Oil, and EU ETS EUA Allowances, and time series split: Training sample (2018–2021), Validation sample (2021–2023), and Test sample (2023–2025).

Summary statistics are presented in Table .4 (in Appendix 7) for the variables during the initial calibration period and the subsequent out-of-sample test period. As we mentioned, electricity prices have become volatile in recent years. In our data, the price range between a minimum of -500 EUR/MWh and a maximum exceeding 930 EUR/MWh in the test window reflects heightened volatility. In addition, the price standard deviation decreased from 112.4 to 50.8 EUR/MWh, indicating more outliers, particularly negative prices.

Figure 2 shows the data split. The initial training window is further subdivided into two subsets: a training set, which consists of the first D_{train} days, and a validation set, which comprises the D_{val} days at the end of the window, two years before the test sample. The split is used exclusively during Optuna’s

³<https://www.lseg.com/en/data-analytics/products/datastream-macroeconomic-analysis>

hyperparameter tuning. As soon as the optimal configuration is identified, it is fixed and reused across all rolling forecasts.

We also adopt a rolling forecasting scheme that mimics daily forecasting operations to ensure a realistic evaluation. As illustrated in Figure 2, for each day of the out-of-sample period, the model is retrained using the most recent $T=1456$ days. Each iteration advances the rolling window by one day, generating a 24-hour forecast for the following day. The window then shifts forward by 1 day, and the process repeats. This continues sequentially until all test days (730 days) are forecasted, which runs from 15 January 2023 to 13 January 2025.

Moreover, parameter warm-starting is one of the key features of the forecasting pipeline: the model parameters obtained at day T are used to initialise the model at day $T + 1$. By recursively reusing weights, training can be significantly stabilised, and convergence times are reduced. During the first step of the rolling loop, this warm-starting is combined with OLS-based initialisation of the linear component for models that contain skip-connections, i.e., linear expert model as one of the components or when it is used as stand-alone model. As a result of this rolling and retraining strategy, along with the separation of training, validation, and test data, a robust assessment of generalisation performance can be achieved, mimicking the operational conditions of market forecasters.




Figure 2: Rolling window and training strategy for electricity price forecasting. The approach consists of three phases: (1) Initial training on 4 years of historical data (from 10 October 2018 to 13 January 2021), (2) Rolling validation with weight transfer between windows (from 14 January 2021 to 14 January 2023), and (3) Rolling testing (from 15 January 2023 to 13 January 2025) with continuous model updates.

4. Methodology

The objective of this study is to develop and evaluate a novel neural network architecture for day-ahead electricity price forecasting. The proposed architecture integrates RNN, Kalman filter and skip connections, aiming to improve predictive accuracy by leveraging both linear and nonlinear modeling capabilities. The proposed approach is used to model the vector of daily electricity prices $\mathbf{Y}_t = (Y_{t,0}, \dots, Y_{t,S-1})$ for $S = 24$ as a dimension. Table 1 summarises the model architectures used in this study, including the linear expert model, the single RNN branch types models, and the different combined architectures.

4.1. Models' Input Features

Our input features are selected to reflect both key economic drivers and electricity market dynamics in order to accurately forecast day-ahead electricity prices.

Most predictors are shared between linear and RNN-based forecasting models. These predictors include historical price data, market expectations, fuel fundamentals, and calendar effects.

1. Lagged electricity prices

Autoregressive behaviour is accounted for by including historical day-ahead prices.

For the linear expert model, we use lags of one, two, and seven days: Y_{t-1} , Y_{t-2} , and Y_{t-7} .

The RNN models, however, includes only Y_{t-1} , since the sequential structure of the network takes into account additional temporal dependencies.

2. Day-ahead fundamental indicators, including:

- Load forecast: Expected demand for the target day, denoted Load_{t+1} .
- Renewable energy forecasts: Onshore and offshore wind generation forecast, denoted Wind_{t+1} . Solar PV forecast, denoted Solar_{t+1} . These reflect the anticipated supply from variable renewable sources.

3. Fuel prices: These serve as proxies for marginal generation costs. Brent oil, denoted Oil_{t-2} . Natural gas, denoted NGas_{t-2} . Hard coal, denoted Coal_{t-2} .

4. Carbon pricing: EU Emission Allowances (EUA) are included to account for CO₂ costs borne by fossil fuel generators, denoted EUA_{t-1} .

5. Calendar dummies: Indicators for specific days of the week (especially Monday, Saturday, and Sunday) are included to capture weekly seasonality, usually weekday dummies, resp. one-hot encoded information and denoted collectively as cal_t .

All features are available before the market closes at noon and reflect the information set accessible to market participants when submitting bids. For linear models, 24 separate outputs are generated for each hour of the next day, based on a fixed input vector. As a result, each output is modelled independently, and the structure makes no allowance for sequential time dynamics beyond the explicit inclusion of lag variables.

RNN inputs consist of multivariate sequences over the past T days, where the sequence length is considered a hyperparameter tunable within the interval $[1, 7]$. This temporal structure enables the network to learn

dynamic patterns such as seasonality or delayed effects in the data. Consequently, lag features beyond Y_t are omitted, as the architecture captures these dependencies implicitly.

As a result, we have two sets of input features:

- Input features for linear expert models

$$\mathbf{X}_{t,s} = (\mathbf{Ylag}_{t,s}, \mathbf{cal}_t, \mathbf{Fund}_{t,s}, \mathbf{price}_t) \quad (1)$$

- Input features for RNN models

$$\mathbf{X}_t^{\text{RNN}} = (\mathbf{Y}_t, \mathbf{cal}_t, \mathbf{Fund}_t, \mathbf{price}_t) \quad (2)$$

4.2. State-of-the-art linear structures

The linear model for all $s = 0, \dots, S - 1$ is given by the following equation:

$$Y_{t+1,s} = \beta_{s,0} + \text{LM}(\mathbf{X}_{t,s}; \beta_s) + \varepsilon_{t+1,s} \quad (3)$$

with $\text{LM}(\mathbf{X}; \beta) = \mathbf{X}\beta$ representing a linear model. In the literature, many variants of linear models are usually referred to as expert models, [Billé et al. \(2023\)](#); [Maciejowska et al. \(2024\)](#); [Uniejewski & Ziel \(2025\)](#). All expert models have the structure $\mathbf{X}_{t,s} = (\mathbf{Ylag}_{t,s}, \mathbf{cal}_t, \mathbf{Fund}_{t,s}, \mathbf{price}_t)$.

Note that, usually, as the notation suggests, $\mathbf{cal}_t$ and $\mathbf{price}_t$ is daily information, while $\mathbf{Ylag}_{t,s}$ and $\mathbf{Fund}_{t,s}$ are high-frequent (hourly) data. Moreover, $\mathbf{cal}_t$ is deterministic information whereas $\mathbf{price}_t$ and $\mathbf{Fund}_{t,s}$ are stochastics.

Here, we utilize the expert model specification from [Ziel \(2016\)](#):

$$\mathbf{Ylag}_{t,s} = (Y_{t-1,s}, Y_{t-2,s}, Y_{t-7,s}) \quad (4)$$

$$\mathbf{cal}_t = (\text{Mon}_{t+1}, \text{Sat}_{t+1}, \text{Sun}_{t+1}) \in \mathbb{R}^{D_{\text{cal}}} \quad (5)$$

$$\mathbf{Fund}_{t,s} = (\text{Load}_{t+1,s}, \text{Wind}_{t+1,s}, \text{Solar}_{t+1,s}) \in \mathbb{R}^{D_{\text{fund}}} \quad (6)$$

$$\mathbf{price}_t = (\text{EUA}_{t-1,s}, \text{NGas}_{t-1,s}, \text{Oil}_{t-1,s}, \text{Coal}_{t-1,s}) \in \mathbb{R}^{D_{\text{price}}} \quad (7)$$

The calendar information $\mathbf{cal}_t$ can be regarded as a one-hot encoded information of the univariate day-of-week time series, taking values from 1 to 7 characterising the corresponding weekday of time t .

This may be written in reduced multivariate form as

$$\mathbf{Y}_t = \beta_0 + \text{LEM}(\mathbf{X}_t, \mathbf{B}) + \varepsilon_t \quad (8)$$

with $\mathbf{B}$ and $\mathbf{X}_t$ as corresponding regression matrix which leads to equivalence to (3).

4.3. The Recurrent Neural Network RNN Model

The RNN ([Elman \(1990\)](#)) is a Feed-forward Neural Network in which the information is transferred from the input layer to the output layer. However, the RNN saves the output of a specific layer and connects back to the input to predict the output. More precisely, the RNN uses their internal state

(memory) to process sequences of inputs with variable length. It consist of three layers: input, hidden, and output layer. Here we consider Elman regression Network given by the following equation:

$$\begin{aligned} \mathbf{H}_t &= \text{RNN}_{\text{cell}}(\mathbf{H}_{t-1}, \mathbf{X}_t^{\text{RNN}}) \\ \mathbf{Y}_t &= \mathbf{W}_{\text{out}} \mathbf{H}_t + \mathbf{b}_{\text{out}} \end{aligned} \quad (9)$$

where $\text{RNN}_{\text{cell}}(\mathbf{H}, \mathbf{X}^{\text{RNN}}; \mathbf{W}_{\text{hid}}, \mathbf{W}_{\text{ext}}, \mathbf{b}_{\text{hid}}) = g(\mathbf{W}_{\text{hid}} \mathbf{H} + \mathbf{W}_{\text{ext}} \mathbf{X}^{\text{RNN}} + \mathbf{b}_{\text{hid}})$ and $\mathbf{X}_{t \in \mathbb{R}^D}^{\text{RNN}}$ where $D = S + D_{\text{cal}} + SD_{\text{fund}} + D_{\text{price}}$, and $\mathbf{Fund}_t = (\text{fund}_{t,0}, \dots, \text{fund}_{t,S-1})$

The hidden state at each time step t is denoted as $\mathbf{H}_t \in \mathbb{R}^H$, where H represents the number of hidden units in the network. The hidden state is updated recursively using the previous hidden state $\mathbf{H}_{t-1} \in \mathbb{R}^H$, the current input vector $\mathbf{X}_{t,s}^{\text{RNN}}$, and the associated weight matrices.

As shown in Figure 3, the transition from the previous hidden state to the current state is governed by the recurrent weight matrix $\mathbf{W}_{\text{hid}} \in \mathbb{R}^{H \times H}$, which captures temporal dependencies, and the input weight matrix $\mathbf{W}_{\text{ext}} \in \mathbb{R}^{h \times D}$, which maps the input features into the hidden state space. A bias term $\mathbf{b}_{\text{hid}} \in \mathbb{R}^h$ is also included to adjust the transformation.




Figure 3: Graph of a simple Elman network with one input, one hidden layer, and one output. The hidden state $\mathbf{h}_t$ is computed from the previous state $\mathbf{h}_{t-1}$ (dashed loop arrow) and the current input $\mathbf{x}_t$. The output $\mathbf{y}_t \in \mathbb{R}^S$ is then obtained through a linear transformation of the hidden state via the output weight matrix $\mathbf{W}_{\text{out}} \in \mathbb{R}^{S \times h}$ and bias term $\mathbf{b}_{\text{out}} \in \mathbb{R}^S$.

The object $\mathbf{g} = (g_1, \dots, g_H)$ represents the vector of activation function. We utilize here the ReLU (Rectified Linear Unit) activation function for g_h to update the hidden state in each time step of the RNN. Where: $\mathbf{g}_h = \mathbf{g}_{\text{ReLU}}(z) = \text{ReLU}(z) = \max(0, z) = z \mathbb{1}_{\{z > 0\}}$, $z \in \mathbb{R}$, $\mathbf{g}_{\text{ReLU}}(z) \in [0, \infty)$.

The output of the RNN at each time step, $\mathbf{Y}_t \in \mathbb{R}^S$, is computed using the hidden state transformation through an output weight matrix $\mathbf{W}_{\text{out}} \in \mathbb{R}^{S \times h}$ and a bias term $\mathbf{b}_{\text{out}} \in \mathbb{R}^S$ (see Figure 3).

4.4. The Kalman Filter as a Special Case of RNN

The Kalman filter [Kalman \(1960\)](#) is a recursive algorithm that estimates the state of a dynamic system from a series of noisy measurements. It operates in two steps: prediction and update. In the context of RNNs, the Kalman filter can be viewed as a special case where the hidden state is updated based on the previous state and the current observation, similar to the recurrent connections in RNNs.

The Kalman filter can be expressed as follows:

$$\begin{aligned} \mathbf{H}_t &= \mathbf{A}_{\text{hid}} \mathbf{H}_{t-1}^{KF} + \mathbf{A}_{\text{ext}} \mathbf{Y}_t + \mathbf{b}_{\text{hid}} \\ \mathbf{Y}_t &= \mathbf{A}_{\text{out}} \mathbf{H}_t^{KF} + \mathbf{b}_{\text{out}} \end{aligned} \quad (10)$$

where $\mathbf{H}_t^{KF}$ is the hidden state at time t , $\mathbf{Y}_t$ is the observation, and $\mathbf{A}_{\text{hid}}$, $\mathbf{A}_{\text{ext}}$, and $\mathbf{A}_{\text{out}}$ are weight matrices. The Kalman filter assumes a linear relationship between the hidden state and the observation, making it suitable for linear systems. Considering identity function $\mathbf{g}_h = \mathbf{g}_{\text{id}}(z) = \text{id}(z) = z$, $z \in \mathbb{R}$, $\mathbf{g}_{\text{id}}(z) \in \mathbb{R}$ as an activation function, instead of ReLU in the case of RNN model, the Kalman filter can be considered as a linear RNN with a specific structure.

Recently, a new sequence modeling architecture called *Mamba* has been introduced as a general form of the state-space model (Gu & Dao (2024)). As an extension of the classical Kalman filter, the Mamba architecture provides flexible sequence modeling capabilities. In a similar way to the Kalman filter, Mamba models time-based dependencies through a hidden state that changes over time. In contrast to the Kalman filter which uses fixed matrices to define state transitions and observations, *Mamba* replaces them with learnable operators that change with input data.

The *Mamba* model is expressed as:

$$\begin{aligned}\mathbf{H}_t^{MB} &= \mathbf{D}_t \mathbf{H}_{t-1}^{MB} + \mathbf{E}_t \mathbf{Y}_t, \\ \mathbf{Y}_t &= \mathbf{C}_t \mathbf{H}_t^{MB},\end{aligned}\tag{11}$$

where $\mathbf{D}_t$, $\mathbf{E}_t$, and $\mathbf{C}_t$ are dynamically computed from the input sequence data. When these matrices are fixed across time, such that: $\mathbf{D}_t = \mathbf{A}_{\text{hid}}$, $\mathbf{E}_t = \mathbf{A}_{\text{ext}}$, $\mathbf{C}_t = \mathbf{A}_{\text{out}}$, the model reduces to the classical linear state-space equations of the Kalman filter as expressed in Equation 10.

Based on this relationship, the Kalman filter can be viewed as a specific case within the larger Mamba architecture, which uses fixed, linear, and time-invariant matrices. Mamba extends this by allowing adaptive, input-dependent dynamics.

4.5. The Parallel-Branch RNNs with Skip Connections

The parallel-branch RNN architecture with skip connections is designed to enhance the model's ability to capture complex temporal dependencies and non-linear relationships in the data. As shown in Figure 4, this architecture consists of multiple parallel branches, each processing the input sequence independently and then merging their outputs. The skip connections allow for direct information flow between layers, facilitating the learning of long-range dependencies.

The combined model will include the non-linear layer by the RNN model, linear layer with time dependencies consideration (Kalman filter, KF) and the linear model (LEM), skip connection. Therefore, the new model will combine Equation 9, Equation 10, and 3. The architecture can be represented as follows:

$$\begin{aligned}\mathbf{Y}_{t+1} &= \beta_0 + \text{LEM}(\mathbf{X}_t, \beta) + \mathbf{W}_{\text{out}} \mathbf{H}_t^{\text{RNN}} + \mathbf{A}_{\text{out}} \mathbf{H}_t^{\text{KF}} + \varepsilon_{t+1} \\ \mathbf{H}_t^{\text{RNN}} &= \text{RNN}_{\text{cell}}(\mathbf{H}_{t-1}^{\text{RNN}}, \mathbf{X}_{t,s}^{\text{RNN}}; \mathbf{W}_{\text{hid}}, \mathbf{W}_{\text{ext}}, \mathbf{b}_{\text{hid}}^{\text{RNN}}) \\ \mathbf{H}_t^{\text{KF}} &= \mathbf{A}_{\text{hid}} \mathbf{H}_{t-1}^{\text{KF}} + \mathbf{A}_{\text{ext}} \mathbf{X}_{t,s}^{\text{RNN}} + \mathbf{b}_{\text{hid}}^{\text{KF}}\end{aligned}\tag{12}$$

LEM input is represented by the vector of exogenous variables $\mathbf{X}_{t,s}$ defined as indicated in Equation 1, and the RNN input and the Kalman filter input are represented by the vector of exogenous variables $\mathbf{X}_{t,s}^{\text{RNN}}$ defined as indicated in Equation 2.




Figure 4: A Parallel-Branch Recurrent Neural Network architecture with skip connections for day-ahead electricity price forecasting. The linear model for direct effects, an RNN branch for nonlinear temporal patterns, and the KF branch represents the Kalman Filter, which is for state-space dynamics. The outputs from the three branches are summed to produce 24-hour-ahead forecasts $(y_1, \dots, y_{24})$.

4.6. Different Model Architectures

We implement and compare the following model architectures, ranging from single models to all possible hybrid and parallel combinations. Table 1 provides a brief explanation of each model architecture, explaining its structure, input features, and how linear and nonlinear components are integrated to capture both the temporal and fundamental dynamics of the market.

- **Stand-alone Models:**

- Elman RNN (with ReLU): Recurrent neural network models of the Elman type, with ReLU activation as non-linear deep neural networks model (see Equation 9).

- Elman RNN (with Identity) or Kalman filter: Recurrent neural network models of the Elman type, with identity activation, with is considered as a linear Kalman-type recurrence model (see Equation 10).
- Linear expert model with OLS-initialized: Baseline linear regression model using all input features and lags, with optional initialization from the ordinary least squares (OLS) solution (see , Equation 3).

- **Hybrid Architectures:**

- Models with skip connection that combine LEM and RNN branches (Equation 9 and Equation 3).
- Models with parallel RNN branches (RNN and Kalman) without skip connections (Equation 9 and Equation 10).
- Models with parallel LEM and RNN branches (RNN and Kalman) with skip connections (Equation 9, Equation 10, and Equation 3).

Type	Abbrev.	Architecture	Description	Input Features
1	RNN	Elman RNN (ReLU)	No skip connections; purely nonlinear processing path.	RNN Input: $\mathbf{X}_t^{\text{RNN}}$
2	KF	Elman RNN (Identity)	RNN with identity activation representing a Kalman filter model.	RNN Input: $\mathbf{X}_t^{\text{RNN}}$
3	LEM	Linear expert model (OLS init)	Purely linear model; can be initialized via ordinary least squares (OLS).	Linear Input: $\mathbf{X}_t$
4	LEM-RNN	Hybrid: Linear + RNN (ReLU, Skip, OLS)	Combined model with a ReLU RNN branch and a linear skip connection. The final output is the sum of both branches.	RNN Branch: $\mathbf{X}_t^{\text{RNN}}$ Linear Skip: $\mathbf{X}_t$
5	KF-RNN	Dual Parallel RNNs (ReLU + Identity)	Two parallel RNNs: RNN & KF. Their outputs are concatenated and passed to a final linear layer.	RNN Inputs: $\mathbf{X}_t^{\text{RNN}}$
6	LEM-KF-RNN	ReLU + Identity + Linear, OLS	Allowing both nonlinear and linear components to contribute to the forecast.	Dual RNN: $\mathbf{X}_t^{\text{RNN}}$ Linear Skip: $\mathbf{X}_t$

Table 1: Overview of implemented model architectures, their structures, and input features.

5. Training and Forecasting Study, and Evaluation Metrics

5.1. Data preprocessing

Data preprocessing ensures that the forecasting framework has appropriate inputs, including temporal consistency and feature completeness. A dataset of hourly day-ahead power prices along with relevant explanatory variables is available for Germany. First, to accommodate daylight saving time (DST) shifts,

a proprietary transformation aligned all hourly data to a consistent temporal grid, ensuring uniform spacing and comparability between seasons. Second, the dataset was restructured into a daily-by-hour matrix representation, with each row corresponding to a calendar day and each column to a specific hour of the day. As a result, a regressor matrix was built by aligning the endogenous variable with its associated regressors. We removed the rows with missing values that resulted from lag construction.

The matrix formed here will serve as the basis for training the model. Finally, we standardise all variables using the mean/standard deviation to ensure comparable scales of contribution from all variables. During each iteration of the rolling-window evaluation, train-only standardisation is applied to both input features and target variables, i.e., the mean (μ_{train}) and (σ_{train}) are calculated based only on the available data on the rolling window, which is important to ensure no data leakage. Through the standardisation process, each variable is transformed so that it has a zero mean and unit variance as follows:

$$\mathbf{x}_i^{\text{std}} = \frac{\mathbf{x}_i - \mu_{\text{train}}}{\sigma_{\text{train}}}, \quad \mathbf{y}_i^{\text{std}} = \frac{\mathbf{y}_i - \mu_{\text{train}}^{(y)}}{\sigma_{\text{train}}^{(y)}} \quad (13)$$

Where $\mathbf{x}_i^{\text{std}}$ represents the standardised feature matrix and $\mathbf{y}^{\text{std}}$ denotes the target value vector. It is necessary to perform this standardisation step to optimise model stability and convergence during training. In order to return predictions to their original units (e.g., €/MWh), the standardised forecasts of the target value, associated with the model outputs, are de-standardised. The inverse transformation is presented as follows:

$$\hat{\mathbf{y}}_i = \hat{\mathbf{y}}_i^{\text{std}} \cdot \sigma_{\text{train}}^{(y)} + \mu_{\text{train}}^{(y)} \quad (14)$$

This method guarantees that evaluation metrics like RMSE, MAE, etc., are calculated in their original scale, facilitating a significant and authentic evaluation of forecasting precision in genuine out-of-sample scenarios.

5.2. Optimization of the objective functions

In our forecasting methodology, we use a dual-level optimization approach with an inner learning process for estimating model parameters and an outer optimisation process for determining the optimal configuration of hyperparameters. In the following subsections, each objective function and its associated estimation process are described in detail.

5.2.1. Estimation and Update of the models parameters

As shown in Equation 12, a number of parameters will be optimised. Depending on the model type, including recurrent weights that connect hidden states over time, such as hidden to hidden connections ($\mathbf{W}_{\text{hid}}$ and $\mathbf{A}_{\text{hid}}$ for RNN and KF, respectively), hidden to output weights ($\mathbf{W}_{\text{out}}$ and $\mathbf{A}_{\text{out}}$), as well as bias terms. Weights from our linear expert model, initialised using OLS estimates, will be also optimised. In order to estimate the neural component parameters, supervised gradient-based learning is employed within a rolling window approach, while LEM coefficients are initialised with OLS estimates rather than learned directly from gradient updates. The model generates forecasts $\hat{\mathbf{Y}}_t^{\text{std}} = f(\mathbf{X}_t^{\text{std}}, \theta)$ for a given set of standardised inputs $\mathbf{X}_t^{\text{std}}$ and target outputs $\mathbf{Y}_t^{\text{std}}$, where θ denotes the set of all trainable parameters (weights and biases) in the network.

Over the training window, the parameters are estimated by minimising the following loss function:

$$\mathcal{L}_{\text{train}}(\theta) = \underbrace{\frac{1}{N} \sum_{t=1}^N \|\hat{y}_t - y_t\|_2^2}_{\text{Mean Squared Error (MSE)}} + \underbrace{\lambda_1 \sum_{i,j} |w_{ij}|}_{\text{L1 Regularization (Sparsity)}} + \underbrace{\lambda_2 \sum_{i,j} w_{ij}^2}_{\text{L2 Regularization (Weight Decay)}}.$$

where L1 regularisation (LASSO) penalises only output-side linear mappings, promoting sparsity and interoperability, while L2 regularisation (Ridge) penalises overfitting by implementing weight decay in the optimiser.

The parameters are updated using the Adam optimisation algorithm (Adam et al. (2014)), which is an adaptive variation of stochastic gradient descent that estimates the gradients' first and second moments to dynamically update learning rates.

At each iteration t , the parameters are updated as follows:

$$m_t = \gamma_1 m_{t-1} + (1 - \gamma_1) g_t, \quad v_t = \gamma_2 v_{t-1} + (1 - \gamma_2) g_t^2, \quad \theta_{t+1} = \theta_t - \eta \frac{m_t}{\sqrt{v_t} + \epsilon}.$$

Where $\mathbf{g}_t = \nabla_{\theta_t} \mathcal{L}_{\text{train}}$ denotes the gradient of the loss, and η is the learning rate. We implement ReduceLROnPlateau scheduling, which reduces the learning rate automatically (by a half) when the validation loss stagnates for five epochs. In addition, we adopt a gradient clipping of $\|\mathbf{g}_t\|_2 \leq 5$ which minimises the risk of gradients exploding in recurrent components.

As the training proceeds, rolling windows are simulated, mimicking real-time learning. Models are re-trained each forecast day based on fixed-length windows of recent observations. Model types differ in terms of how the weights corresponding to the initial window (D_{init}) are initialised, with RNN models starting with small uniform random weights and LEM models starting with OLS-based warm-starts.

Once convergence has been achieved for the current window, the optimised weights are saved as initialisation weights for the next window (D_{all}). While maintaining learned temporal dependencies, this sequential re-estimation enables the parameters to change smoothly across time. We thus implement a continuously adaptable model in the estimation process by combining serial parameter transfer between windows with stochastic gradient-based learning within each window.

5.2.2. Optimization and Hyperparameter Training

The previous section explains how the model's (internal) weights (θ) are estimated and updated on a rolling basis. The following section will discuss hyperparameter optimisation, which is an outer optimisation procedure that optimises a set of hyperparameters (ϕ), responsible for the model's architecture and learning dynamics.

An overview of the adopted key hyperparameters is explained in Table 2. These hyperparameters are tuned using `optuna` (Akiba et al. (2019)) with the Tree-structured Parzen Estimator (TPE) sampler for Bayesian optimization, the models are implemented in PyTorch (Paszke et al. (2019)), which efficiently explores the search space. We also adopt a rolling window forecasting approach (as described in Section 3) where a search space is designed to strike a balance between computing efficiency, convergence stability, and model adaptability.

For a given `optuna` trial k , the entire model training and validation cycle is evaluated based on a specific

Table 2: Overview of tuned hyperparameters and their search spaces used in Optuna (TPE) optimization.

Name	Symbol	Type / Interval	Description
Hidden layer size	H	$\{1, 2, \dots, 128\}$	Number of neurons in each RNN branch.
Sequence length	L	$\{1, 2, \dots, 7\}$	Number of past days fed into the RNN.
Initial training window size	D_{init}	$[30, 730]$	Number of days used for the first rolling forecast window.
Updated training window size	D_{all}	$[2, 365]$	Number of days used for subsequent rolling windows.
Initial epochs	E_{init}	$\{10, 20, 50, 100\}$	Number of training epochs for the initial window.
Updated epochs	E_{all}	$\{5, 10, 20, 50\}$	Number of training epochs for each updated rolling step.
Initial learning rate	η_{init}	$[10^{-5}, 10^{-2}]$	Adam learning rate for the initial forecast.
Updated learning rate	η_{all}	$[10^{-4}, 10^{-2}]$	Adam learning rate for rolling updates.
Initial weight decay (L_2)	$\lambda_{w,\text{init}}$	$[10^{-8}, 10^{-2}]$	L_2 regularization for the initial window.
Updated weight decay (L_2)	$\lambda_{w,\text{all}}$	$[10^{-8}, 10^{-2}]$	L_2 regularization for subsequent windows.
Initial L_1 regularization	$\lambda_{1,\text{init}}$	$[10^{-6}, 10^{-1}]$	L_1 penalty applied in the initial window.
Updated L_1 regularization	$\lambda_{1,\text{all}}$	$[10^{-6}, 10^{-1}]$	L_1 penalty applied in rolling updates.
OLS initialization weight	α	$[0, 2]$	Scaling factor for OLS-based weight initialization.
OLS initialization flag	<code>use_ols_weights</code>	$\{\text{True}, \text{False}\}$	Enables warm-start from OLS weights (skip/linear models).
Batch size	B	$\{8, 16, 32, 64\}$	Mini-batch size used during optimization.
Gradient clip norm	c_{grad}	$[0.1, 10]$	Maximum gradient norm for clipping (stability).
Dropout rate	p_{drop}	$[0.0, 0.5]$	Dropout applied to hidden layers for regularization.

hyperparameter configuration ϕ_k . A trial’s performance is determined by its Root Mean Square Error (RMSE), which is calculated over all forecast hours and evaluation days as follows:

$$\mathcal{L}_{\text{train}}(\phi_k) = \text{RMSE}_{\text{overall}} = \sqrt{\frac{1}{ST} \sum_{s=1}^S \sum_{t=1}^T (\hat{y}_{s,t} - y_{s,t})^2}.$$

where S denotes the number of hours of the day $S = 0, 1, \dots, 23$, and T represents the number of days in the evaluation period (rolling window).

The set of tuned hyperparameters ϕ consist of:

The hyperparameters of the RNN architecture are hidden layer size (H) and sequence length (L). RNNs with larger hidden dimensions can capture complex nonlinear dynamics. Sequence length determines how many past observations are available to learn temporal dependencies (ranging here from one day to one week).

The size of the training dataset is determined by the data-window parameters D_{init} and D_{all} in initial and subsequent rolling estimation. This allows our models to take into consideration the structural changes in the underlying data-generating process as additional information is made available, ensuring that it evolves adaptively as new information becomes available.

We consider two *training schedule parameters*: E_{init} and E_{all} , which specify how many optimisation epochs are allocated to each window. The initial window has a greater number of epochs to ensure stable convergence, while successive windows have fewer training to accelerate updates as previously learned weights are refined.

Concerning the *optimisation hyperparameters*, we adopt the learning rate of the Adam optimiser η_{init} and η_{all} , which scale the gradient $\nabla_{\theta}\mathcal{L}$ during parameter updates, controlling the step size in the optimisation space. The optimiser can progress more quickly towards a minimum with a high learning rate (η) value, but it risks skipping it, which may result in unstable training or even divergence. A slower training process may arise from a lower learning rate, even while it guarantees smoother, more stable convergence.

In addition, to improve generalisation and control model complexity, our loss function integrates both Ridge ($L2$) and Lasso ($L1$) regularisation terms. The result is an Elastic Net penalty that balances shrinkage and sparsity as follows: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MSE}} + \lambda_w \|\mathbf{W}\|_2^2 + \lambda_1 \|\mathbf{W}\|_1$, Where MSE loss function is defined as follows: $\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$, $\|\mathbf{W}\|_2^2 = \sum_j w_j^2$ denotes the weight decay (ridge) penalty, and $\|\mathbf{W}\|_1 = \sum_j |w_j|$ represents the sparsity of the L1 penalty. The coefficients λ_w and λ_1 control the strength of each term separately for initial and rolling windows.

For models with linear or skip connections (model types 3, 4, and 6), the linear branch is warm started using coefficients estimated from OLS. OLS based initialisations are controlled by the scaling factor α and the Boolean flag `use_ols_weights`. Note that: $\mathbf{W}_{\text{OLS}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$, and $\mathbf{W}_{\text{init}} = \alpha \cdot \mathbf{W}_{\text{OLS}}$, where $\alpha \in [0, 2]$ adjusts the relative contributions from the OLS solution to the initial weight configuration, which facilitates faster convergence and ensures stable results.

There are *three training control parameters*: batch size B , gradient clipping norm c_{grad} , dropout probability p_{drop} , and activation type σ . Euclidean gradients are constrained by gradient clipping, satisfying: $\|\nabla_{\theta}\mathcal{L}\|_2 \leq c_{\text{grad}}$, which improves recurrent architectures by mitigating exploding gradients. Dropout deactivates hidden neurons randomly with probability p_{drop} , resulting in enhanced generalisation due to a reduction in co-adaptation.

The objective is to minimize the RMSE. Each `optuna` trial runs for 500 iterations, evaluating different hyperparameter configurations. The best-performing set is then selected. In the rolling window training approach, the best hyperparameters found in the first window using `optuna` are stored and reused as starting values for subsequent windows instead of being randomly selected (See Figure 2). Firstly, the model is initialized with random weights (or with OLS weights), but after each iteration, the best hyperparameters are saved and loaded into the next estimation window. In addition to ensuring continuity in learning, this approach accelerates training and improves stability by retaining previously learned information. As the model evolves rather than starting from scratch, it effectively captures evolving patterns in our data.

5.2.3. Out-of-sample rolling window forecasting

After optimising the weights of the models (θ^*) and the hyperparameter configuration (ϕ^*), the next step consists of computing out-of-sample forecasts using ϕ^* and θ^* . Our test sample consists of the last two years of the dataset (730 days). During testing, the rolling-window procedure is followed, ensuring that each forecast is produced using only the most recent information available. Based on the most recent historical data, the model receives standardised input features $\mathbf{X}_t^{\text{std}}$, and produces standardised predictions $\hat{\mathbf{Y}}_t^{\text{std}} = f(\mathbf{X}_t^{\text{std}}, \theta^*)$. These forecasts are then de-standardised (as shown in Equation 14) to be transformed back to their original scale and compared to the corresponding real values in the test

sample.

Hybrid architectures using LEM and/or multiple RNN components, such as model types 4–6, compute the final forecast as the sum of several standardised sub-forecasts associated with the model’s internal components. Let $\hat{Y}_t^{std}(c)$ be the standardised forecast produced by component c (i.e, LEM, ReLU-RNN, or identity-RNN branches, i.e., KF branch). After combining these components, we obtain the total standardised forecast: $\hat{Y}_t^{std} = \sum_c \hat{Y}_t^{std}(c)$. To restore forecasts at their original scale, the aggregated standardised output is de-standardised based on the mean and standard deviation determined from the corresponding training window:

$$\hat{Y}_t = \left(\sum_c \hat{Y}_t^{std}(c) \right) \sigma_y^{(train)} + \mu_y^{(train)}. \quad (15)$$

By analysing the forecasts of stand-alone models, we can quantify the contributions of linear and nonlinear components to the overall forecast, which enables a transparent evaluation of the contributions of each sub-model to the final forecast.

5.2.4. Forecasting evaluation

To evaluate the forecasting performance of the proposed models, a variety of evaluation metrics are used, in addition, we assess their statistical significance.

Evaluation metrics:

We utilise two widely accepted metrics: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE). This metric quantifies the average magnitude of errors between predicted and actual electricity prices, providing insight into the accuracy of the forecast. Furthermore, the Naive weekly model is used to calculate the relative Mean Absolute Error (rMAE).

The model predicts electricity prices for the next 24 hours, generating an output vector: $\hat{\mathbf{Y}}_t = (\hat{Y}_{t,0}, \hat{Y}_{t,1}, \dots, \hat{Y}_{t,S-1})$.

To evaluate the accuracy of the model in the test sample, forecasted prices $\hat{\mathbf{Y}}_t$ are compared with the actual prices $\mathbf{Y}_t$ from the test set. Then RMSE and MAE over the test period are calculated as follow:

$$\text{RMSE} = \sqrt{\frac{1}{ST} \sum_{t=1}^T \sum_{s=0}^{S-1} (Y_{t,s} - \hat{Y}_{t,s})^2}, \quad \text{MAE} = \frac{1}{ST} \sum_{t=1}^T \sum_{s=0}^{S-1} |Y_{t,s} - \hat{Y}_{t,s}| \quad (16)$$

Where T is the number of days in the test sample, $\hat{Y}_{t,s}$ is the forecasted price for day t at hour s and $Y_{t,s}$ is the actual price (real price value). As well as the RMSE and the MAE for every individual hour:

$$\text{RMSE}_s = \sqrt{\frac{1}{T} \sum_{t=1}^T (Y_{t,s} - \hat{Y}_{t,s})^2}, \quad \text{MAE}_s = \frac{1}{T} \sum_{t=1}^T |\hat{Y}_{t,s} - Y_{t,s}| \quad (17)$$

The RMSE and MAE provide insights into the model’s performance, with lower values indicating better accuracy.

The relative MAE (rMAE) (Hyndman & Koehler (2006)) normalises the error by the MAE obtained from a naive forecasting model. Likewise, naive forecasts are constructed using out-of-sample data. In this study, we applied the weekly Naive model, defined as follows:

$$\hat{Y}_{t,s}^{\text{naive}} = Y_{t-7,s} \quad (18)$$

This choice is motivated by the fact that it is easy to compute, and unlike the rMAE based on random walk model ($\hat{Y}_{t,s}^{\text{naive},1} = Y_{t-1,s}$), it captures weekly effects.

Hence, the rMAE is described as:

$$\text{rMAE} = \frac{\text{MAE}}{\text{MAE}_{\text{naive}}}$$

Evaluation Based on Statistical Testing.

The analysis of evaluation metrics must also be accompanied by consideration of whether any difference in accuracy is statistically significant. In other words, to determine whether forecast accuracy differences are real or just random differences between forecasts, we need to perform statistical tests. Diebold-Mariano (DM) test (Diebold & Mariano (2002)) and Giacomini-White (GW) test (Giacomini & White (2006)) have been widely used to test the statistical significance of electricity price forecasting. However, the GW test is considered preferable as it can be viewed as a generalisation of the DM test (Lago et al. (2021), López et al. (2025), and Marcjasz et al. (2020)).

In our research, a multivariate GW test is performed jointly for all hours using the multivariate loss differential series or the daily loss differential series defined as:

$$\Delta_d^{\text{A,B}} = \|\varepsilon_d^{\text{A}}\|_p - \|\varepsilon_d^{\text{B}}\|_p, \quad (19)$$

where $\|\cdot\|_1$ denotes the L_1 norm. For each day, the error vector ε_d is 24-dimensional, with: $\varepsilon_d = (\varepsilon_{d,1}, \varepsilon_{d,2}, \dots, \varepsilon_{d,24})$ and $\varepsilon_{d,h} = P_{d,h} - \hat{P}_{d,h}$, where $P_{d,h}$ is the observed value and $\hat{P}_{d,h}$ the forecasted value. Based on Equation 19, the GW test yields a single p -value assessing the null hypothesis:

$$H_0 : \mathbb{E}(\Delta_{A,B}) \leq 0 \quad (20)$$

which indicates that model A has lower forecasting error than model B. The alternative hypothesis is:

$$H_A : \mathbb{E}(\Delta_{A,B}) \geq 0 \quad (21)$$

6. Empirical Results and discussion

Empirical results, including hyperparameter optimisation outcomes and forecasting performance evaluation, are presented and discussed in this section.

6.1. Hyperparameter optimization results

Figure 5 captures the variation in the influence of relevant variables on electricity prices throughout the day based on the linear estimation component (LEM) estimations from Model LEM-RNN-KF. The highest impact is attributed to the last price of the previous day. It shows a high influence weight in the early hours and decreases in the evening hours. The lagged prices maintain a moderate positive weight, suggesting the temporal persistence of prices is strong with temporal persistence, while the installed loads have a steady positive influence all day. Wind and solar generation show very slight and even negative influence weights, thus suggesting a price-reducing impact. Similarly, the variable solar generation maintains a position of consistently negative weight over the day, especially between 10 am and 15 am. This aligns with standard expectations of the price-suppressing impact of solar generation

during daylight hours. The fuel and carbon price variables, namely oil, gas, coal, and EUA, have influence weights that are close to zero, suggesting a weak influence in the short run.

Table .5 in Appendix summarises hyperparameter optimisation results from the validation sample for




Figure 5: Features importance from the LEM component of the LEM-KF-RNN model across all hours.

all model architectures. This set of optimal hyperparameters will serve as initial values for forecasting the test sample. Despite the pure RNN achieving the lowest validation RMSE (28.12), the difference between it and the best hybrid configuration, especially the RNN-LEM-KF model (28.16), is marginal (≈ 0.05 RMSE). The low RMSE value associated with the RNN model can be the result of finely tuned hyperparameters. More precisely, the optimal initial learning rate is small ($\eta_{\text{init}} = 2.48 \times 10^{-5}$), the initial training window is large ($D_{\text{init}} = 359$), and mild regularisation ($\lambda_{w,\text{init}} = 5.31 \times 10^{-4}$) that preserves nonlinear temporal dependencies while stabilising recurrent gradients. The linear LEM and KF models, conversely, depend on explicit feature relationships and limited memory, as evidenced by their learning rates ($\eta_{\text{init}} = 1.71 \times 10^{-5}$ and 1.23×10^{-3} , respectively), abbreviated update intervals (e.g., $D_{\text{all}} = 33$ for KF model), and more stringent penalties (e.g., $\lambda_{w,\text{init}} = 3.09 \times 10^{-3}$ for KF model). This explains their relatively low validation accuracy (RMSE = 31.30 and 30.52, respectively).

The optimal sequence length of all models is set to 1 ($L = 1$), signifying that extended sequences were unnecessary. Furthermore, across models with an LEM component, the consistent selection of the OLS-based initialisation indicates that this is crucial to enhance interpretability and convergence stability. However, the relative ranking of models during the test phase may not necessarily hold even though the RNN has the lowest numerical validation RMSE.

To assess the importance of hyperparameters in the learning process, we performed a fANOVA analysis (Hutter et al. (2014)). Results in Figure 6 indicate that the length of the rolling window is the most powerful hyperparameter (≈ 0.74), highlighting the importance of historical data in enhancing model accuracy. Followed by learning rate, which, with moderate significance (≈ 0.20), is essential for achieving steady convergence. Other hyperparameters showed similar minor effects. This indicates that historical




Figure 6: Important variables identified during hyperparameter optimization.

data and learning rate play a greater role in forecasting accuracy than other architectural or regularisation adjustments; the former plays a primary role, and the latter plays a secondary role.

Figure 7 illustrates the impact of each hyperparameter (initialisation and rolling-update values) on




Figure 7: Top hyperparameter values for best-performing trials.

the model's performance (RMSE) during 500 iterations of the training period. For the learning rate, both the initial (η_{init}) and updated values (η_{all}) perform better in the interval $[10^{-4}, 10^{-3}]$, meaning in mid-range values. However, the RMSE increased when learning speeds were extremely high or low, which indicates that the learning rate ought to be balanced. In the case of weight decay ($\lambda_{w,init}$ and $\lambda_{w,all}$), a mild regularisation of about 10^{-4} and 10^{-3} range tends to reduce overfitting. Thus, a moderate L2 regularisation value is beneficial during initialisation and rolling adaptation. Concerning the L1 regularization parameters ($\lambda_{l,init}$ and $\lambda_{l,all}$), a moderate value of L1 regularization assists the model in

focusing on salient features rather than irrelevant ones. In terms of sequence length, models with short input sequences of one to three time steps or days yielded the best performance (lower RMSE), indicating that long historical sequences do not provide additional useful patterns. Finally, the rolling window size (D_{all}) significantly improved accuracy over longer training periods, especially between 300 and 365 days, implying that the model is capable of understanding more complex patterns over longer periods.

These findings are reinforced by the results in Figure .13 in Appendix, which illustrates how each hyperparameter was explored during training and where the search trials were most concentrated. It indicates that the updated training window D_{all} is clearly left-skewed (most trials are packed near the high end ~ 330 –365 with a tail towards small windows). Both L1 regularisation parameters peaked near 0 ($\approx 10^{-3}$ – 10^{-4}) with a tail for large values. This is also similar for the learning rate and the weight decay. A left-skewed distribution of neurons, with a strong tendency towards larger sizes (100–128), suggests that a model with larger hidden layers is able to capture more complex patterns and nonlinear relationships in the data, thus improving forecasting. The sequence length is strongly right-skewed (over 400 trials chose 1). α is left-skewed (mass near 1.5–2.0), indicating that stronger OLS scaling was favoured. In addition, use of OLS weights are dominated by True.

Figure 8 illustrated that the lowest RMSE was achieved early during the optimization process, suggest-




Figure 8: Hyperparameter optimization progress across 500 trials, where the blue dots represent the RMSE values achieved for each trial, the new best RMSE is indicated by the green dots, the red line links best RMSE values over trials.

ing that effective hyperparameter values were found rather quickly, while the true global minimum was found much later, around trial 400. The randomly scattered results of the trials and the lack of an overall downward trend provide evidence for the stochasticity of the search process, and imply that further trials, beyond 500, could provide marginal returns.

6.2. Forecasting results

An interpretation and discussion of the evaluation metrics, the GW statistical test, and the forecast decomposition analysis are provided in this section.

6.2.1. Forecasting accuracy metrics

A summary of forecasting accuracy metrics for all models over the two-year testing period is illustrated in Table 3.

Model	RMSE	MAE	rMAE
Weekly Naive	51.617	32.998	1.00
RNN	23.078	13.446	0.407
Kalman Filter	25.352	14.059	0.426
LEM	24.465	13.719	0.415
LEM-RNN	24.940	13.579	0.411
Kalman Filter-RNN	23.503	13.226	0.400
LEM-Kalman Filter-RNN	22.754	13.006	0.394
LEAR	26.587	16.098	0.487
DNN	25.976	13.725	0.415

Table 3: Forecasting accuracy metrics for different models. Subtle color gradients highlight performance, with **greener shades indicating lower (better) errors**, and **redder shades indicating higher (worse) errors**

For the single-type models, we noticed that the RNN model outperformed the stand-alone linear models, with an RMSE of 23.078. This indicates its strength in capturing the nonlinear dynamics and time dependences that characterise electricity price time series. However, purely linear models, i.e., LEM and KF, generate higher forecasting errors (RMSE = 24.465 and RMSE = 25.352, respectively). Despite their ability to capture structural trends and mean-reverting patterns, linear methods often struggle with nonlinearity, spikes, and volatility.

For the combined models, we found that the hybrid models significantly improved the forecasting accuracy of the linear components (LEM and KF) by incorporating a nonlinear component (RNN). Specifically, KF-RNN achieved lower errors (RMSE = 23.503) compared to LEM-RNN (RMSE = 24.940). Furthermore, the LEM-KF-RNN model achieved the best forecasting results overall (RMSE = 22.754) by fully integrating linear regression (LEM), state-space filtering (KF), and nonlinear sequence learning (RNN). This strong performance comes from the model’s ability to capture both smooth structural components through LEM and KF, and rapid, nonlinear changes through RNN. Combining these elements allows the model to respond well to different market conditions and keep a good balance between bias and variance.

To assess the accuracy of the proposed model compared to other models in the literature, we first

consider the naive model (see Equation 18). Furthermore, we benchmark our models with the state-of-the-art models proposed by Lago et al. (2021). This includes:

1. A parameter-rich autoregressive model with exogenous variables, known as the Lasso Estimated Autoregressive (LEAR) model. Estimation of the proposed LEAR model is based on L1-regularisation (LASSO, Least Absolute Shrinkage and Selection Operator).
2. A Deep Neural Network (DNN) model, which extends the traditional multilayer perceptron (MLP). This network consists of four layers, uses a multivariate framework (single model with 24 outputs), is estimated using Adam Adam et al. (2014), and its hyperparameters and input features are optimised using a Bayesian optimisation algorithm called the tree Parzen estimator Bergstra et al. (2011).

As input features, both models use the same set of information:

- a. Electricity price lagged values (historical day-ahead prices of the prior three days p_{d-1} , p_{d-2} , p_{d-3} , and one week ago p_{d-7}).
- b. A day-ahead load forecast $\mathbf{x}_d^1$ along with the previous day $\mathbf{x}_{d-1}^1$ and one week ago $\mathbf{x}_{d-7}^1$.
- c. Similarly, the forecasted day-ahead solar and wind generation $\mathbf{x}_d^2$ plus the previous day value $\mathbf{x}_{d-1}^2$ and the value a week ago $\mathbf{x}_{d-7}^2$.
- d. The weekday is represented as a dummy variable.

Detailed information about the two state-of-the-art models can be found in Lago et al. (2021). The python codes are open access and available in Github⁴.

Examining the benchmark models, the Weekly Naive baseline performs as expected, with the highest errors (RMSE = 51.617), since it does not account for structural or temporal variability beyond simple weekly repetition. We also examined the advanced LEAR and DNN models. Even though they are sophisticated, they still have relatively low accuracy (RMSE = 26.587 and 25.976, respectively). A possible explanation for this is that fuel prices are not taken into account in the set of input features despite their direct effects on electricity prices (Ghelasi & Ziel (2024)), especially during the energy crisis of 2022 (see Figure 1).

In summary, These findings highlight the benefits of hybrid models combining linear, filtering, and non-linear components for accurate forecasting. This can be achieved through puzzling out various features in electricity prices through each model's strengths.

6.2.2. The multivariate GW test results

According to the results of the multivariate GW test using the L_1 norm in Equation 19, we have the following loss differential series:

$$\Delta_d^{A,B} = \sum_{h=1}^{24} |\varepsilon_{d,h}^A| - \sum_{h=1}^{24} |\varepsilon_{d,h}^B| \quad (22)$$

To illustrate the range of obtained p -values, we use heat maps arranged as chessboards. Given that this test is run at a 5% significance level, results are interpreted based on p -values as follows:

⁴<https://github.com/jeslago/epftoolbox>

- $p - \text{values} < 0.05$ (green) indicate significant outperformance of a model on the X-axis (better) compared to the model on the Y-axis (worse).
- $p - \text{values} > 0.05$ indicate significant underperformance of the model on the X-axis (worse) compared to the model on the Y-axis (better).




Figure 9: Giacomini-White test results for all models over the two-year test period. For each pair of models, two one-sided tests are run, and a heat map is used to show the range of p-values. The closer they are to zero (dark green), the greater the difference between the forecasts of models on the X-axis (better) and those on the Y-axis (worse). A black indicates a p-value greater than 10%.

The GW results are illustrated in Figure 9, the following findings were gathered:

- As indicated by the green cells in the RNN column, RNN outperforms KF and LEM, which is consistent with its lower RMSE and MAE values. It also significantly outperforms the LEAR and DNN state of the art models.
- Both KF and LEM columns are in black, meaning that all other models consistently outperform them, and their forecasting accuracy is not statistically significant, while there is no statistical significance between their forecasting accuracy. Additionally, these models significantly outperform the benchmark models (LEAR and DNN), except for the DNN and KF models, where the latter does not significantly outperform the former.

- The combination models (LEM-RNN, KF-RNN, and LEM-KF-RNN) outperformed all single model components and all benchmark models, proving the effectiveness of hybridisation. There is an exception, however, in the RNN and LEM-RNN pair, which does not show a significant difference at 5%, suggesting that LEM does not substantially improve RNN.
- It is worth mentioning that the differences between the hybrid models themselves are mostly statistically insignificant (black cells), meaning that the LEM-KF-RNN, KF-RNN, and LEM-RNN are all comparable in forecasting accuracy when compared to each other.

In summary, the performance of hybrid models is clearly superior to that of single models. The best architectures, KF-RNN and LEM-KF-RNN, show similar statistical results. Consequently, the most effective combinations appear to have reached a plateau in performance.

Figure 10 visualises the hourly RMSE during the test period. It indicates that the LEM-KF-RNN model replicates the daily price cycle accurately, demonstrating a precise understanding of both magnitude and phase as it shows lower prices during the morning, an increase at noon, and a significant peak in the evening. More precisely, the figure shows a varying RMSE value throughout the day, peaking in the late afternoon at about 35, when peak demand occurs. During the night hours, however, a lower value (10–15) is observed (base hours).




Figure 10: RMSE by Hour - Hourly Test Set.

6.2.3. Decomposed forecasts

A comprehensive view of the behaviour of the hybrid LEM-KF-RNN model and its comparison to the real electricity price and different forecast components for two years out-of-sample testing period is presented in Figures 11 and 12. The Figures illustrate how the combined forecast gains the advantage from its linear and nonlinear components in enhancing forecasting accuracy on both a daily and hourly basis.

To better visualize the individual forecasts alongside the combined forecast and actual prices, each component was unstandardized separately following Equation (14).

Let $\hat{\mathbf{Y}}_t^{\text{std}}(c)$ denote the standardized forecast of component c ($c = \text{LEM}, \text{KF}, \text{RNN}$); then the corresponding unstandardized forecast $\hat{\mathbf{Y}}_t(c)$ is obtained as:

$$\hat{\mathbf{Y}}_t(c) = \hat{\mathbf{Y}}_t^{\text{std}}(c)\sigma_y + \mu_y \quad (23)$$

It is important to note that the de-standardized combined forecast (Equation 15) is not the direct sum of the de-standardized components, since the summation is performed on the standardized values. Consequently, the combined forecast time series (shown in red in Figures 11 and 12) does not represent a simple sum of the individual unstandardized forecasts.

First, Figure 11a illustrates the difference in average daily forecasts for stand-alone and combined models. It illustrates the fact that the combined forecast (red line) closely tracks the actual price (dashed line), capturing both long-term trends and short-term fluctuations. As for the model's components, we noticed that the LEM component (yellow) can model baseline structural movements, which is consistent with its ability to track smooth, low-frequency variations in electricity prices. Alternatively, the RNN components (ReLU and identity activations) capture high-frequency deviations such as peaks, drops, and volatility spikes, with the RNN model with ReLU showing greater superiority over the KF component. Therefore, combining these different components results in a model that is strong and interacts well, enabling the aggregated output to track persistent trends and transient shocks, thus predicting prices with high accuracy regardless of market conditions.




Figure 11: Decomposed, Combined & Real Electricity Prices for the Test Period. Subfigure (a) shows the **real price values** (dashed blue line) compared with forecasted values from the **Combined LEM-KF-RNN model**, which represents the aggregated forecast from all model branches: **RNN with ReLU forecast** and **RNN with Identity forecast** correspond to the two RNN branches (ReLU and linear activations), and **the LEM model component**. Subfigure (b) presents the same decomposition at the hourly level for the test set. All values are unstandardized.

Secondly, Figure 11b shows forecasts by hour of the day, confirming that the combined model follows the intraday demand-price cycle accurately. There is a great deal of agreement between the predicted and actual hourly patterns, especially during peak hours in the afternoon and evening, when volatility is high. The intraday price cycle was also captured by the model's components with variation in amplitude and accuracy. The RNN model was found to be more accurate (closest to the actual and combined value)

followed by the KF model and then the LEM model.

Lastly, Figure 12 examine component contributions during the first and last test days. In the first test day plot (Figure 12a) for Sunday, 15 January 2023, real electricity prices remained relatively stable and very low. This may have been the result of a calm market on the weekend or an increased share of renewables, which reduces price volatility. In spite of this, the linear component still produces a typical daily cycle with pronounced peak and base hours, thus predicting variation that does not occur in reality. Alternatively, the RNN branches respond more adaptively: the identity-based branch exhibits inverse movement of the LEM model, thus cancelling the overestimation of the latter and reducing artificial peaks introduced by it, a process further supported by the ReLU-based branch. By combining these adjustments, the combined model can generate a more realistic forecast that reflects the dynamics of a low-volatility day. Therefore, the combined forecast seems close to that day’s stable market behaviour. According to the last test day plot (Figure 12b) (Monday, 13 January 2025), prices show clear intraday variability, with peaks in early morning and late evening, which will be typical of winter demand on a working day. In general, the linear model captures the shape of these peaks, but underestimates their intensity. Alternatively, RNNs seem more capable of forecasting the amplitude of price peaks and troughs, resulting in a forecast that’s more accurate at capturing short-term volatility and follows actual price dynamics.

In summary, the decomposition plots illustrate how different components contribute to the combined model depending on market conditions. Although LEM models can impose typical daily patterns even on stable days, RNN and KF components respond adaptively, suppressing unnecessary cycles when the market is calm and emphasising them when strong intraday dynamics appear. This adaptive equilibrium renders the hybrid forecast more realistic and more aligned with actual price fluctuations.




Figure 12: Decomposed Forecasts for the First and Last Day of the Test Set.

7. Summary and Conclusion

Electricity price forecasting has been a challenging task due to the complex nature of the underlying electricity markets. A wide variety of EPF models have been studied in the past, including statistical models and machine learning techniques, as well as hybrid frameworks that combine both of type models. In this context, our paper introduced a new hybrid framework that combines a linear expert model

(LEM) with RNNs and KF using a parallel-branch neural network architecture. Hence, the proposed architecture benefits from the strength of linear and nonlinear methods, making it possible to model complicated market interdependencies in time and structure. Our methodology is empirically tested on the day-ahead German electricity market using hourly data covering the period between 2018 and 2025. we also considered a set of input features that reflect the main drivers of electricity prices. to assess the accuracy of our model we benchmark it with state -of-the-art models from the literature. Results demonstrate the improvement in forecasting accuracy of the combined architectures compared to stand-alone models. From the stand-alone model results, the RNN model outperformed the linear LEM and KF models, emphasising the importance of nonlinearity dynamics estimation in electricity price. Concerning the hybrid models, our findings demonstrated that the best model performance was realised by the LEM-KF-RNN hybrid model, closely followed by the KF-RNN hybrid model. This emphasizes the importance of incorporating time sequence learning into both nonlinear and linear model components.

Concerning the decomposition analysis, although the RNN components are shown to be dynamically changing with highly volatile market conditions, variable relationships can be explained using the LEM branch by providing interpretability. Moreover, the KF branch enhances linear temporal dependencies, affirming the complementary function of state-space modelling within the hybrid model.

These results align with the findings from [Lago et al. \(2021\)](#) and [Amor et al. \(2024\)](#) that support the outperformance of deep learning models over statistical models for electricity price forecasting. Our findings also align with a lot of research which developed hybrid or combined models (e.g., [Ben Amor et al. \(2018\)](#), [Zhang et al. \(2020\)](#), etc.), who demonstrated that combining linear and non-linear models improves forecasting accuracy compared to both single components. Another advantage of such hybrid model types is discussed by [Rondón-Cordero et al. \(2025\)](#) in the context of energy consumption forecasting.

Our methodology can be further developed in future research. First, investigating more advanced state-space models, such as the Mamba model [Gu & Dao \(2024\)](#), as this model allows more expressive and adaptive (dynamic) temporal dependencies. Secondly, rather than employing a linear summing of the component forecasts in the neural network output layer, the model could dynamically and adaptively ascertain the contribution of each branch, identify the most effective components based on forecast accuracy, and assign appropriate weights ([Xu et al. \(2024\)](#) and [Palou et al. \(2025\)](#)). Finally, since probabilistic forecasting is important in decision-making related to electricity trading and system operations, the proposed model architecture may be extended to probabilistic forecasting ([Ziel & Steinert \(2018\)](#) and [Nowotarski & Weron \(2018\)](#)).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was partially funded in the course of TRR 391 Spatio-temporal Statistics for the Transition of Energy and Transport (520388526) by the Deutsche Forschungsgemeinschaft (DFG, German

Research Foundation).

Appendix

Table .4: Descriptive statistics of the data used for the initial calibration window and the out-of-sample test period.

Series	Mean	Std	Min	Q25	Median	Q75	Max
In-sample data (01.10.2018–14.01.2023)							
Price (EUR/MWh)	97.2	112.4	-90.0	34.0	50.1	112.4	871.0
Load forecast (MWh)	55.2	9.2	32.3	47.7	55.2	63.1	77.6
Wind forecast (MWh)	14.2	10.4	0.2	5.9	11.4	20.3	47.7
PV forecast (MWh)	5.2	8.1	0.0	0.0	0.2	7.9	37.8
EUA (EUR/tCO ₂)	44.7	24.6	15.2	24.6	29.3	67.8	97.6
Coal (EUR/t)	117.5	100.4	34.7	49.1	64.4	147.2	402.7
Gas (EUR/MWh)	49.4	56.1	3.5	12.8	19.8	83.7	310.5
Brent Oil (EUR/bbl.)	62.3	21.4	17.8	51.5	58.3	73.0	117.5
Out-of-sample data (15.01.2023–13.01.2025)							
Price (EUR/MWh)	86.7	50.8	-500.0	63.0	88.8	112.1	936.3
Load forecast (MWh)	52.5	9.0	30.5	45.1	52.2	59.9	73.3
Wind forecast (MWh)	15.9	11.5	0.2	6.4	13.1	23.3	50.3
PV forecast (MWh)	6.8	10.4	0.0	0.0	0.2	10.9	48.2
EUA (EUR/tCO ₂)	74.4	10.8	50.5	66.3	72.3	84.5	97.1
Coal (EUR/t)	109.9	12.5	86.6	102.0	108.5	114.7	162.1
Gas (EUR/MWh)	37.6	8.4	23.0	31.2	36.4	43.2	67.7
Brent Oil (EUR/bbl.)	74.9	5.5	62.7	70.5	74.8	78.5	91.9

Table .5: Comparison of tuned hyperparameters across all model types (Optuna, TPE).

Symbol	RNN	LEM	KF	LEM-RNN	KF-RNN	RNN-LEM-KF
H	107	–	17	88	108	118
L	1	–	1	1	1	1
D_{init}	512	712	512	97	648	302
D_{all}	359	65	33	358	357	358
η_{init}	2.48×10^{-5}	1.71×10^{-5}	1.23×10^{-3}	1.00×10^{-5}	2.84×10^{-3}	8.05×10^{-4}
η_{all}	4.85×10^{-4}	6.81×10^{-4}	1.57×10^{-4}	3.65×10^{-4}	4.80×10^{-4}	1.09×10^{-4}
$\lambda_{w,\text{init}}$	5.31×10^{-4}	3.77×10^{-4}	3.09×10^{-3}	4.85×10^{-3}	1.38×10^{-3}	9.48×10^{-4}
$\lambda_{w,\text{all}}$	4.40×10^{-4}	4.27×10^{-4}	1.05×10^{-4}	1.30×10^{-3}	4.97×10^{-4}	1.67×10^{-3}
$\lambda_{1,\text{init}}$	1.01×10^{-5}	3.58×10^{-3}	1.24×10^{-3}	1.72×10^{-4}	2.01×10^{-4}	1.41×10^{-4}
$\lambda_{1,\text{all}}$	5.56×10^{-4}	1.65×10^{-3}	2.31×10^{-3}	4.34×10^{-4}	1.06×10^{-3}	1.79×10^{-4}
α	–	0.756	–	1.000	–	1.113
use_ols_weights	–	True	–	True	–	True
Best RMSE	28.118	31.300	30.515	28.164	28.308	28.162




Figure .13: Hyperparameter Sampling Distributions.

References

- Abedinia, O., Amjady, N., Shafie-Khah, M., & Catalão, J. P. (2015). Electricity price forecast using combinatorial neural network trained by a new stochastic search method. *Energy Conversion and Management*, 105, 642–654. doi:<https://doi.org/10.1016/j.enconman.2015.08.025>.
- Adam, K. D. B. J. et al. (2014). A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 1412.
- Aggarwal, S. K., Saini, L. M., & Kumar, A. (2009). Electricity price forecasting in deregulated markets: A review and evaluation. *International Journal of Electrical Power & Energy Systems*, 31, 13–22.
- Ahmad, W., Javaid, N., Chand, A., Shah, S. Y. R., Yasin, U., Khan, M., & Syeda, A. (2019). Electricity price forecasting in smart grid: A novel e-cnn model. In *Workshops of the International Conference on Advanced Information Networking and Applications* (pp. 1132–1144). Springer.
- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 2623–2631).
- Amor, S. B., Möbius, T., & Müsgens, F. (2024). Bridging an energy system model with an ensemble deep-learning approach for electricity price forecasting. *arXiv preprint arXiv:2411.04880*, .

- Anbazhagan, S., & Kumarappan, N. (2014). Day-ahead deregulated electricity market price forecasting using neural network input featured by dct. *Energy conversion and management*, 78, 711–719. doi:<https://doi.org/10.1016/j.enconman.2013.11.031>.
- Baldassi, C., Malatesta, E. M., Perugini, G., & Zecchina, R. (2023). Typical and atypical solutions in nonconvex neural networks with discrete and continuous weights. *Physical Review E*, 108, 024310.
- Bejani, M. M., & Ghatee, M. (2021). A systematic review on overfitting control in shallow and deep neural networks. *Artificial Intelligence Review*, 54, 6391–6438.
- Ben Amor, S., Akkal Devi, P., & Müsgens, F. (2024). Meta-forecasting for solar power generation: algorithm-based swarm intelligence. *20th International Conference on the European Energy Market (EEM)*, .
- Ben Amor, S., Boubaker, H., & Belkacem, L. (2018). Forecasting electricity spot price for nord pool market with a hybrid k-factor gamma-llwnn model. *Journal of Forecasting*, 37, 832–851. doi:<https://doi.org/10.1002/for.2544>.
- Bento, P., Pombo, J., Calado, M., & Mariano, S. (2018). A bat optimized neural network and wavelet transform approach for short-term price forecasting. *Applied energy*, 210, 88–97.
- Bergstra, J., Bardenet, R., Bengio, Y., & Kégl, B. (2011). Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, 24.
- Billé, A. G., Gianfreda, A., Del Grosso, F., & Ravazzolo, F. (2023). Forecasting electricity prices with expert, linear, and nonlinear models. *International Journal of Forecasting*, 39, 570–586.
- Camara, A., Feixing, W., Xiuqin, L. et al. (2016). Energy consumption forecasting using seasonal arima with artificial neural networks models. *International Journal of Business and Management*, 11, 231. doi:[10.5539/ijbm.v11n5p231](https://doi.org/10.5539/ijbm.v11n5p231).
- Castello, O., & Resta, M. (2025). Univariate and multivariate forecasting of the electricity futures curve using dynamic recurrent neural networks. *Applied Energy*, 394, 126082.
- de Castro Matias, M., & Tabak, B. M. (2025). Comparison of indicator saturation and markov regime-switching models for brazilian electricity prices. *Energy Economics*, 144, 108341.
- Chai, S., Li, Q., Abedin, M. Z., & Lucey, B. M. (2024). Forecasting electricity prices from the state-of-the-art modeling technology and the price determinant perspectives. *Research in International Business and Finance*, 67, 102132.
- Chan, S.-C., Tsui, K. M., Wu, H., Hou, Y., Wu, Y.-C., & Wu, F. F. (2012). Load/price forecasting and managing demand response for smart grids: Methodologies and challenges. *IEEE signal processing magazine*, 29, 68–85. doi:[10.1109/MSP.2012.2186531](https://doi.org/10.1109/MSP.2012.2186531).
- Chen, Z., Zhang, B., Du, C., Yang, C., & Gui, W. (2025). Outlier-adaptive-based non-crossing quantiles method for day-ahead electricity price forecasting. *Applied Energy*, 382, 125328.
- Chinnathambi, R. A., Plathottam, S. J., Hossen, T., Nair, A. S., & Ranganathan, P. (2018). Deep neural networks (dnn) for day-ahead electricity price markets. In *2018 IEEE electrical power and energy conference (EPEC)* (pp. 1–6). IEEE.
- Chu, F.-L. (2009). Forecasting tourism demand with arma-based methods. *Tourism Management*, 30, 740–751. doi:[http://dx.doi.org/10.1016/j.tourman.2008.10.016](https://dx.doi.org/10.1016/j.tourman.2008.10.016).
- Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. (2014). Empirical evaluation of gated recurrent neural networks on sequence modeling. arxiv 2014. *arXiv preprint arXiv:1412.3555*, 1412.
- Darudi, A., Bashari, M., & Javidi, M. H. (2015). Electricity price forecasting using a new data fusion algorithm. *IET Generation, Transmission & Distribution*, 9, 1382–1390.
- Dautel, A. J., Härdle, W. K., Lessmann, S., & Seow, H.-V. (2020). Forex exchange rate forecasting using deep recurrent neural networks. *Digital Finance*, 2, 69–96.

- Diebold, F. X., & Mariano, R. S. (2002). Comparing predictive accuracy. *Journal of Business & economic statistics*, 20, 134–144.
- Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14, 179–211.
- Garcia, R. C., Contreras, J., Van Akkeren, M., & Garcia, J. B. C. (2005). A garch forecasting model to predict day-ahead electricity prices. *IEEE transactions on power systems*, 20, 867–874. doi:[10.1109/TPWRS.2005.846044](https://doi.org/10.1109/TPWRS.2005.846044).
- Ghelasi, P., & Ziel, F. (2024). From day-ahead to mid and long-term horizons with econometric electricity price forecasting models. *arXiv preprint arXiv:2406.00326*, .
- Ghelasi, P., & Ziel, F. (2025). From day-ahead to mid and long-term horizons with econometric electricity price forecasting models. *Renewable and Sustainable Energy Reviews*, 217, 115684.
- Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74, 1545–1578.
- Gianfreda, A., Ravazzolo, F., & Rossini, L. (2020). Comparing the forecasting performances of linear models for electricity prices with high res penetration. *International Journal of Forecasting*, 36, 974–986.
- Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, .
- Gu, A., & Dao, T. (2024). Mamba: Linear-time sequence modeling with selective state spaces. In *First conference on language modeling*.
- Hao, Y., Tian, C., & Wu, C. (2020). Modelling of carbon price in two real carbon trading markets. *Journal of Cleaner Production*, 244, 118556. doi:<https://doi.org/10.1016/j.jclepro.2019.118556>.
- Hewamalage, H., Ackermann, K., & Bergmeir, C. (2023). Forecast evaluation for data scientists: common pitfalls and best practices. *Data Mining and Knowledge Discovery*, 37, 788–832.
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2, 359–366.
- Hung, Y.-C., Lin, F.-J., Hwang, J.-C., Chang, J.-K., & Ruan, K.-C. (2014). Wavelet fuzzy neural network with asymmetric membership function controller for electric power steering system via improved differential evolution. *IEEE Transactions on Power Electronics*, 30, 2350–2362. doi:[10.1109/TPEL.2014.2327693](https://doi.org/10.1109/TPEL.2014.2327693).
- Hussain, S., Teni, A. P., Hussain, I., Hussain, Z., Pallonetto, F., Eichman, J., Irshad, R. R., Alwayle, I. M., Alharby, M., Hussain, M. A. et al. (2024). Enhancing electric vehicle charging efficiency at the aggregator level: A deep-weighted ensemble model for wholesale electricity price forecasting. *Energy*, 308, 132823.
- Hutter, F., Hoos, H., & Leyton-Brown, K. (2014). An efficient approach for assessing hyperparameter importance. In *International conference on machine learning* (pp. 754–762). PMLR.
- Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International journal of forecasting*, 22, 679–688.
- Jiang, P., Yang, H., & Ma, X. (2019). Coal production and consumption analysis, and forecasting of related carbon emission: evidence from china. *Carbon Management*, 10, 189–208. doi:<https://doi.org/10.1080/17583004.2019.1577177>.
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems, .
- Kanamura, T., & Bunn, D. W. (2022). Market making and electricity price formation in japan. *Energy Economics*, 107, 105765.
- Keles, D., Scelle, J., Paraschiv, F., & Fichtner, W. (2016). Extended forecast methods for day-ahead electricity spot prices applying artificial neural networks. *Applied energy*, 162, 218–230. doi:<https://doi.org/10.1016/j.apenergy.2015.09.087>.

- Khashei, M., & Bijari, M. (2010). An artificial neural network (p, d, q) model for timeseries forecasting. *Expert Systems with applications*, 37, 479–489. doi:<https://doi.org/10.1016/j.eswa.2009.05.044>.
- Ko, M.-S., Lee, K., Kim, J.-K., Hong, C. W., Dong, Z. Y., & Hur, K. (2020). Deep concatenated residual network with bidirectional lstm for one-hour-ahead wind power forecasting. *IEEE Transactions on Sustainable Energy*, 12, 1321–1335.
- Krishna Prakash, N., Singh, J. G. et al. (2023). Electricity price forecasting using hybrid deep learned networks. *Journal of Forecasting*, 42, 1750–1771.
- Kumar, V., Singh, N., Singh, D. K., & Mohanty, S. (2017). Short-term electricity price forecasting using hybrid sarima and gjr-garch model. In *Networking Communication and Data Knowledge Engineering: Volume 1* (pp. 299–310). Springer.
- Kuo, P.-H., & Huang, C.-J. (2018). An electricity price forecasting model by hybrid structured deep neural networks. *Sustainability*, 10, 1280. doi:<https://doi.org/10.3390/su10041280>.
- Lago, J., De Ridder, F., & De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221, 386–405.
- Lago, J., Marcjasz, G., De Schutter, B., & Weron, R. (2021). Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark. *Applied Energy*, 293, 116983.
- Lehna, M., Scheller, F., & Herwartz, H. (2022). Forecasting day-ahead electricity prices: A comparison of time series and neural network models taking external regressors into account. *Energy Economics*, 106, 105742.
- Li, Y., Ding, Y., Liu, Y., Yang, T., Wang, P., Wang, J., & Yao, W. (2022). Dense skip attention based deep learning for day-ahead electricity price forecasting. *IEEE Transactions on Power Systems*, 38, 4308–4327.
- Lin, W.-M., Gow, H.-J., & Tsai, M.-T. (2010). An enhanced radial basis function network for short-term electricity price forecasting. *Applied Energy*, 87, 3226–3234. doi:<https://doi.org/10.1016/j.apenergy.2010.04.006>.
- López, M. Z., Ioannou, Y., & Zareipour, H. (2025). Forecasting electricity prices with deep learning and dynamic sparse training. *Sustainable Energy, Grids and Networks*, (p. 101865).
- Lora, A. T., Santos, J. M. R., Expósito, A. G., Ramos, J. L. M., & Santos, J. C. R. (2007). Electricity market price forecasting based on weighted nearest neighbors techniques. *IEEE Transactions on Power Systems*, 22, 1294–1301. doi:[10.1109/TPWRS.2007.901670](https://doi.org/10.1109/TPWRS.2007.901670).
- Luo, H., & Shao, Y. (2024). Advanced optimal system for electricity price forecasting based on hybrid techniques. *Energies* (19961073), 17.
- Maciejowska, K., Serafin, T., & Uniejewski, B. (2024). Probabilistic forecasting with a hybrid factor-qra approach: Application to electricity trading. *Electric Power Systems Research*, 234, 110541.
- Maciejowska, K., & Weron, R. (2015). Short-and mid-term forecasting of baseload electricity prices in the uk: The impact of intra-day price relationships and market fundamentals. *IEEE Transactions on power systems*, 31, 994–1005.
- Marcjasz, G., Lago, J., & Weron, R. (2020). Neural networks in day-ahead electricity price forecasting: single vs. multiple outputs. *arXiv preprint arXiv:2008.08006*, .
- Marcjasz, G., Narajewski, M., Weron, R., & Ziel, F. (2023). Distributional neural networks for electricity price forecasting. *Energy Economics*, 125, 106843.
- Meng, A., Wang, P., Zhai, G., Zeng, C., Chen, S., Yang, X., & Yin, H. (2022). Electricity price forecasting with high penetration of renewable energy using attention-based lstm network trained by crisscross optimization. *Energy*, 254, 124212.
- Misiorek, A., Trueck, S., & Weron, R. (2006). Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics & Econometrics*, 10.

- Mohamed, N., Ahmad, M. H., Ismail, Z., & Suhartono, S. (2010). Short term load forecasting using double seasonal arima model. In *Proceedings of the regional conference on statistical sciences* (pp. 57–73). volume 10.
- Mujeeb, S., Javaid, N., Ilahi, M., Wadud, Z., Ishmanov, F., & Afzal, M. K. (2019). Deep long short-term memory: A new price and load forecasting scheme for big data in smart cities. *Sustainability*, *11*, 987.
- Mulla, S., Pande, C. B., & Singh, S. K. (2024). Times series forecasting of monthly rainfall using seasonal auto regressive integrated moving average with exogenous variables (sarimax) model. *Water Resources Management*, *38*, 1825–1846.
- Munian, P., & Ziel, F. (2020). Probabilistic forecasting in day-ahead electricity markets: Simulating peak and off-peak prices. *International Journal of Forecasting*, *36*, 1193–1210.
- Niu, X., & Wang, J. (2019). A combined model based on data preprocessing strategy and multi-objective optimization algorithm for short-term wind speed forecasting. *Applied Energy*, *241*, 519–539.
- Nogales, F. J., Contreras, J., Conejo, A. J., & Espínola, R. (2002). Forecasting next-day electricity prices by time series models. *IEEE Transactions on power systems*, *17*, 342–348. doi:[10.1109/TPWRS.2002.1007902](https://doi.org/10.1109/TPWRS.2002.1007902).
- Nowotarski, J., Raviv, E., Trück, S., & Weron, R. (2014). An empirical comparison of alternative schemes for combining electricity spot price forecasts. *Energy Economics*, *46*, 395–412. doi:<https://doi.org/10.1016/j.eneco.2014.07.014>.
- Nowotarski, J., & Weron, R. (2018). Recent advances in electricity price forecasting: A review of probabilistic forecasting. *Renewable and Sustainable Energy Reviews*, *81*, 1548–1568.
- O'Connor, C., Bahloul, M., Rossi, R., Prestwich, S., & Visentin, A. (2025). Conformal prediction for electricity price forecasting in the day-ahead and real-time balancing market. *Energy and AI*, *21*, 100571.
- Ortiz, M., Ukar, O., Azevedo, F., & Múgica, A. (2016). Price forecasting and validation in the spanish electricity market using forecasts as input data. *International Journal of Electrical Power & Energy Systems*, *77*, 123–127. doi:<https://doi.org/10.1016/j.ijepes.2015.11.004>.
- Ortner, A., & Totschnig, G. (2019). The future relevance of electricity balancing markets in europe-a 2030 case study. *Energy Strategy Reviews*, *24*, 111–120.
- Palou, J. T. V., González, J. S., Santos, J. M. R., & Fernández, J. M. R. (2025). A novel weight-based ensemble method for emerging energy players: an application to electric vehicle load prediction. *Energy and AI*, *20*, 100510.
- Panapakidis, I. P., & Dagoumas, A. S. (2016). Day-ahead electricity price forecasting via the application of artificial neural network based models. *Applied Energy*, *172*, 132–151. doi:<https://doi.org/10.1016/j.apenergy.2016.03.089>.
- Paraschiv, D. M., Bălăsoiu, N., Ben-Amor, S., & Bag, R. C. (2023). Hybridising neurofuzzy model and the seasonal autoregressive model for electricity price forecasting on germany's spot market. *Amfiteatru Economic*, *25*, 463–478.
- Paschalidou, E. G., & Thomaidis, N. S. (2025). Risk factors in the formulation of day-ahead electricity prices: Evidence from the spanish case. *Energy Economics*, *142*, 108102.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L. et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, *32*.
- Peng, L., Liu, S., Liu, R., & Wang, L. (2018). Effective long short-term memory with differential evolution algorithm for electricity price prediction. *Energy*, *162*, 1301–1314.
- Pesenti, A., & O'Sullivan, A. (2025). Explaining deep neural network models for electricity price forecasting with xai. *Energy and AI*, (p. 100532).
- Pindoriya, N., Singh, S., & Singh, S. (2008). An adaptive wavelet neural network-based energy price forecasting in electricity markets. *IEEE Transactions On power systems*, *23*, 1423–1432. doi:[10.1109/TPWRS.2008.922251](https://doi.org/10.1109/TPWRS.2008.922251).

- Ramos, P., Santos, N., & Rebelo, R. (2015). Performance of state space and arima models for consumer retail sales forecasting. *Robotics and computer-integrated manufacturing*, 34, 151–163. doi:<http://dx.doi.org/10.1016/j.rcim.2014.12.015>.
- Rondón-Cordero, V. H., Montuori, L., Alcázar-Ortega, M., & Siano, P. (2025). Advancements in hybrid and ensemble ml models for energy consumption forecasting: results and challenges of their applications. *Renewable and Sustainable Energy Reviews*, 224, 116095.
- Sandhu, H. S., Fang, L., & Guan, L. (2016). Forecasting day-ahead price spikes for the ontario electricity market. *Electric Power Systems Research*, 141, 450–459. doi:<https://doi.org/10.1016/j.epsr.2016.08.005>.
- Shahidehpour, M., Yamin, H., & Li, Z. (2003). *Market operations in electric power systems: forecasting, scheduling, and risk management*. John Wiley & Sons.
- Singh, N., Hussain, S., & Tiwari, S. (2018). A pso-based ann model for short-term electricity price forecasting. In *Ambient Communications and Computer Systems: RACCCS 2017* (pp. 553–563). Springer.
- Soares, L. J., & Medeiros, M. C. (2008). Modeling and forecasting short-term electricity load: A comparison of methods with an application to brazilian data. *International Journal of Forecasting*, 24, 630–644. doi:<https://doi.org/10.1016/j.ijforecast.2008.08.003>.
- Tan, Z., Zhang, J., Wang, J., & Xu, J. (2010). Day-ahead electricity price forecasting using wavelet transform combined with arima and garch models. *Applied energy*, 87, 3606–3610.
- Tang, L., Yu, L., & He, K. (2014). A novel data-characteristic-driven modeling methodology for nuclear energy consumption forecasting. *Applied Energy*, 128, 1–14. doi:<https://doi.org/10.1016/j.apenergy.2014.04.021>.
- Tschora, L., Pierre, E., Plantervit, M., & Robardet, C. (2022). Electricity price forecasting on the day-ahead market using machine learning. *Applied Energy*, 313, 118752.
- Tselika, K., Tselika, M., & Demetriades, E. (2024). Quantifying the short-term asymmetric effects of renewable energy on the electricity merit-order curve. *Energy Economics*, 132, 107471.
- Tseng, F.-M., Yu, H.-C., & Tzeng, G.-H. (2002). Combining neural network model with seasonal time series arima model. *Technological forecasting and social change*, 69, 71–87. doi:[https://doi.org/10.1016/S0040-1625\(00\)00113-X](https://doi.org/10.1016/S0040-1625(00)00113-X).
- Ugurlu, U., Oksuz, I., & Tas, O. (2018). Electricity price forecasting using recurrent neural networks. *Energies*, 11, 1255.
- Uniejewski, B., Nowotarski, J., & Weron, R. (2016). Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies*, 9, 621.
- Uniejewski, B., & Ziel, F. (2025). Probabilistic forecasts of load, solar and wind for electricity price forecasting. *arXiv preprint arXiv:2501.06180*, .
- Valenzuela, O., Rojas, I., Rojas, F., Pomares, H., Herrera, L. J., Guillén, A., Marquez, L., & Pasadas, M. (2008). Hybridization of intelligent techniques and arima models for time series prediction. *Fuzzy sets and systems*, 159, 821–845. doi:<https://doi.org/10.1016/j.fss.2007.11.003>.
- Wang, D., Luo, H., Grunder, O., Lin, Y., & Guo, H. (2017). Multi-step ahead electricity price forecasting using a hybrid model based on two-layer decomposition technique and bp neural network optimized by firefly algorithm. *Applied Energy*, 190, 390–407. doi:<https://doi.org/10.1016/j.apenergy.2016.12.134>.
- Wang, J., & Wang, J. (2021). A new hybrid forecasting model based on sw-lstm and wavelet packet decomposition: A case study of oil futures prices. *Computational Intelligence and Neuroscience*, 2021, 7653091.
- Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International journal of forecasting*, 30, 1030–1081.

- Weron, R., & Misiorek, A. (2008). Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models. *International journal of forecasting*, 24, 744–763.
- Xiao, L., Shao, W., Yu, M., Ma, J., & Jin, C. (2017). Research and application of a hybrid wavelet neural network model with the improved cuckoo search algorithm for electrical power system forecasting. *Applied energy*, 198, 203–222.
- Xie, X., Xu, W., & Tan, H. (2018). The day-ahead electricity price forecasting based on stacked cnn and lstm. In *International Conference on Intelligent Science and Big Data Engineering* (pp. 216–230). Springer.
- Xiong, H., & Mamon, R. (2019). A higher-order markov chain-modulated model for electricity spot-price dynamics. *Applied Energy*, 233, 495–515.
- Xiong, X., & Qing, G. (2023). A hybrid day-ahead electricity price forecasting framework based on time series. *Energy*, 264, 126099.
- Xu, Y., Li, J., Wang, H., & Du, P. (2024). A novel probabilistic forecasting system based on quantile combination in electricity price. *Computers & Industrial Engineering*, 187, 109834.
- Yan, W., Wang, P., Xu, R., Han, R., Chen, E., Han, Y., & Zhang, X. (2025). A novel mid-and long-term time-series forecasting framework for electricity price based on hierarchical recurrent neural networks. *Journal of the Franklin Institute*, 362, 107590.
- Yan, X., & Chowdhury, N. A. (2014). Mid-term electricity market clearing price forecasting: A multiple svm approach. *International Journal of Electrical Power & Energy Systems*, 58, 206–214. doi:<https://doi.org/10.1016/j.ijepes.2014.01.023>.
- Yang, H., & Schell, K. R. (2024). Attnet: An explainable gated recurrent unit neural network for high frequency electricity price forecasting. *International Journal of Electrical Power & Energy Systems*, 158, 109975.
- Zhang, J., Tan, Z., & Wei, Y. (2020). An adaptive hybrid model for short term electricity price forecasting. *Applied Energy*, 258, 114087. doi:<https://doi.org/10.1016/j.apenergy.2019.114087>.
- Zhang, T., Tang, Z., Wu, J., Du, X., & Chen, K. (2022). Short term electricity price forecasting using a new hybrid model based on two-layer decomposition technique and ensemble learning. *Electric Power Systems Research*, 205, 107762.
- Zhang, X., Wang, J., & Gao, Y. (2019). A hybrid short-term electricity price forecasting framework: Cuckoo search-based feature selection with singular spectrum analysis and svm. *Energy Economics*, 81, 899–913. doi:<https://doi.org/10.1016/j.eneco.2019.05.026>.
- Zhao, W., Wang, J., & Lu, H. (2014). Combining forecasts of electricity consumption in china with time-varying weights updated by a high-order markov chain model. *Omega*, 45, 80–91. doi:<https://doi.org/10.1016/j.omega.2014.01.002>.
- Zhao, Z., Wang, C., Nokleby, M., & Miller, C. J. (2017). Improving short-term electricity price forecasting using day-ahead lmp with arima models. In *2017 IEEE power & energy society general meeting* (pp. 1–5). IEEE.
- Ziel, F. (2016). Forecasting electricity spot prices using lasso: On capturing the autoregressive intraday structure. *IEEE Transactions on Power Systems*, 31, 4977–4987.
- Ziel, F., & Steinert, R. (2018). Probabilistic mid-and long-term electricity price forecasting. *Renewable and Sustainable Energy Reviews*, 94, 251–266.
- Ziel, F., & Weron, R. (2018). Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics*, 70, 396–420.