

Optimizing Energy Price Forecasting with Hybrid Deep Learning

MSc in Data Analytics
Research Project

Shashank Rahul Shetye
Student ID: x23278510

School of Computing
National College of Ireland

Supervisor: Muslim Jameel Syed

National College of Ireland
Project Submission Sheet
School of Computing



Student Name:	Shashank Rahul Shetye
Student ID:	x23278510
Programme:	Msc in Data Analytics
Year:	2024-2025
Module:	Research Project
Supervisor:	Muslim Jameel Syed
Submission Due Date:	15/09/2025
Project Title:	Optimizing Energy Price Forecasting with Hybrid Deep Learning
Word Count:	6158
Page Count:	17

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	Shashank Rahul Shetye
Date:	15th September 2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	<input type="checkbox"/>
Attach a Moodle submission receipt of the online project submission , to each project (including multiple copies).	<input type="checkbox"/>
You must ensure that you retain a HARD COPY of the project , both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Optimizing Energy Price Forecasting with Hybrid Deep Learning

Shashank Rahul Shetye
x23278510

Abstract

The precise forecasting of energy prices is pertinent to the smooth operation of energy markets, strategic considerations, and, hence, the integration of renewable sources in an increasingly volatile energy market. The research presents a new hybrid deep learning architecture for predicting the hourly market electricity price that synergistically combines Convolutional Neural Networks for spatial temporal feature extraction and Bidirectional Long Short-Term Memory networks for coding both forward and backward temporal dependencies into price volatility. The model considers multi-dimensional datasets containing historical electricity price data, energy consumption data, and comprehensive weather factors encompassing temperature, humidity, wind speed, and precipitation. Overall, the methodology discussed in this paper mainly focused on a hybrid approach: the study conduct numerous experimental evaluations where the CNN-BiLSTM model was compared to individual Gated Recurrent Unit (GRU) and LSTM models using the same datasets and metrics. Meanwhile, this work contributes to furthering the energy price forecasting techniques and can give able support to energy traders, market operators, and utility companies aiming at trying robust forecasting solutions to support their own strategic decisions in weather dependent electricity markets.

Keywords: Electricity price forecasting, hybrid deep learning, CNN-BiLSTM, weather data integration, comparative analysis, GRU, LSTM, energy markets, bi-directional processing

1 Introduction

With market volatility increasing and the varying complexity levels of power markets, energy price forecasting, for one, becomes a more critical determination and challenge among various stakeholders including market participants, energy producers, utility companies, regulatory bodies, and policy analysts.(13) Since electricity in itself should serve as the basic constituent for powering various sectors such as the household, industries, agriculture, health, transportation, communication, and essential public utility services, it needs to be forecast correctly. Reliable access to electricity is not just an economic necessity but sustains the very fundament for sustainable growth and environmental cleanliness, which finally reflects in the living standards of society.(17) The accelerated growth of renewable energy sources have added non-linearities in electricity prices confounded with temporal dependencies arising out of several exogenous parameters including weather conditions, adverse impact due to demand fluctuations, and geopolitical trending.(10)




Figure 1: Implementation

Electricity prices depend on various factors including electricity demand, weather conditions (temperature, wind, rain), coal and gas prices, emissions costs, availability of generation units, etc(2). Even when temporarily converted into mechanical energy by pumped storage plants or into chemical energy through battery systems, electric energy simply cannot be economically stored under normal operating conditions.(18) The current constraints of storage technologies are such that electricity must be consumed practically as soon as it has been generated. Thus, supply and demand (which are both very strong functions of time) have a direct influence on price volatility in the wholesale market, which tends to result in sharp short-term price fluctuations.(6)

A major part of this price volatility arises from the intermittency of renewable energy generation, especially from wind and solar sources. Generation from these sources is actually affected by the weather, which may undergo sudden, unpredictable condition changes.(15) This uncertainty finds a direct way into the price mechanism, making it harder to reach a balance point in the electricity market. At certain places in the research paper, this part studies how different categories of energy influence the price dynamics, while also exploring to what extent irregular weather phenomena, such as unanticipated temperature variations, storms, or cloudiness, may impact both supply and demand aspects of the balance. These intertemporal constraints add another layer of complexity to the planning of energy dispatch and to resource allocation.(20) These requirements result in a spontaneous growth of financial instruments such as electricity futures contracts, with the conduct of trading in these contracts used to hedge price risk. Simultaneously, it sets forth the strategic and analytic importance of considering price forecasts with horizons of longer periods than the typical day ahead forecast window to support more accurate decision making in both the physical and financial markets.(7) Deep learning methods offer a promising approach to modeling complex temporal patterns and long-term dependencies of time series data.(3)

The primary objectives are:

- To develop a hybrid deep learning model that captures temporal dependencies and exogenous influences on electricity prices.
- To compare the performances of the proposed model against standalone recurrent models with similar evaluation metrics and data sets.
- Data-driven feature analysis toward identifying those factors with the strongest impact on prices.

In the growing field of energy informatics, this work is considered to have a meaningful contribution due, basically, to the novel forecasting architecture proposed, which is based on a hybrid CNN-biLSTM model within the setting of high-resolution, hourly electricity price prediction.

2 Related Work

By decarbonization, decentralization, there has been a noticeable shift in the global energy sector with the weather of profound changes.(5)

Energy-use patterns are shifting, becoming more coupled with reality, and more dynamic by the day. About the same set of factors that add to complexity are, among others, real-time weather variability; time-varying consumption demand, which manifests through peak and off-peak charging and discharging, respectively; and, perhaps most importantly, evolving ad hoc consumer demands, shaped by smart technologies, motives toward sustainability, and cost-efficiency concerns.(9) All these reasons put forth in a nutshell, show the dire urgency of flexible forecasting models that are data-enabled and can track fast-changing conditions.

In addition to great progresses in machine learning and data analytics, several persistent and critical issues plague electricity price forecasting models. Handling messy, incomplete, or noisy data from real life is one of them, as it usually requires substantial preprocessing and filtering.

A further major limitation lies in the models' failing to adapt to changes in consumption and generation patterns that are increasingly influenced by intermittent renewables, demand-side management, and climate-related anomalies. Interpretability and computational efficiency are still big issues. To put it simply, most state-of-the-art models are too complex for interpretation by domain experts and stakeholders, too static to accommodate dynamic environments, or too computationally expensive to be deployed at scale in operational settings.

Cheng and Li (2023) performed electricity demand forecasting based on many hourly weather variables, such as temperature, humidity, and wind speed./ The model they suggested reduced RMSE values by roughly 22 percent as compared to the traditional ARIMA-based prediction method, showing the importance of considering meteorological data of high resolution.Despite exhibiting improved performance, unless usually regularization or data augmentation is performed on smaller or less diverse datasets, the model had a tendency to overfit.(4)

In similar pursuits, Singh and Zhao (2023) stressed the continued necessity for temporal feature engineering by adding variables that refer to the hour of the day, month of the year, and public holidays to their forecasting models.(12) Their findings showed significant gains in predictive metrics upon the inclusion of temporal variables, strongly indicating the importance of time contextual information in energy forecasting. Taken together, the literature argued that the forward looking variables, in particular weather, are very important; equally important though, is the time structure of the data for enhancing model performance and thereby its reliability in operational settings.

2.1 Impact of Weather on Energy Demand and Pricing

Weather conditions do basically everything to shape energy demand and price. Temperature, wind speed, humidity, solar radiation, and cloud coverage have been widely accepted to rank among the top predictors in energy forecast models.(11) This is even more prominent for systems with an increasing number of renewable energy sources. Since wind and solar power generation and supply are extremely weather dependent and intermittent, changes in weather sometimes mean major and abrupt changes in the amount of energy supplied at any given time. When paired with variations in demand triggered by loads

dependent on those very variations in weather, the resultant combined effects see direct and sometimes volatile influences on electricity market prices.

In this context, Kumar et al. (2023) carried out an empirical assessment concerning the influence of wind regimes on the hourly electricity prices in the Spanish power market. Their results elucidated that there is a strong and consistent correlation between wind speed variability and supply side price movements, especially in regions where wind power constitutes a significant share in the generation portfolio.(1)

Fernandez and Yu (2024) took a somewhat different approach and factored in real-time meteorological data into their forecasting framework, including humidity, and cloud cover. The findings showed that higher humidity levels were frequently followed by abrupt surges in electricity demand, especially for southern areas where air conditioning is widespread and often imperative during hot days.(8) That is, they accurately captured these precipitous surges, thereby stressing that local weather elements play a crucial role in predicting short-term demand.

Together, these studies underscore the importance of incorporating weather data into energy demand and pricing models. They also reflect a common theme in the literature: while weather-aware forecasting improves model accuracy, challenges remain in automating these systems, adapting them to shifting conditions, and customizing them to specific regional contexts.

2.2 Economic Forecasting and Pricing Models

The issue of using weather informed energy consumption forecasts to calculate monthly energy bills or to design different pricing strategies has been looked into by several other researchers, with studies aimed at making energy costs more responsive to real-time demand patterns for the benefit of both providers and consumers.

Gonzalez and Ahmed (2023) had a forecasting methodology that estimated monthly electricity bills by aggregated weather and load features at sub-daily time steps. By combining climatic variables with past consumption trends, they concluded that a past customer behavior- among which the preceding month's bill as a feature stood out greatly could be a very strong predictor of future energy bills.(14)

While many works provide robust methods for forecasting either energy load, price, or consumption, the following limitations persist:

- Few consider both weather and energy together in a practical billing context.
- Overlook real-world variables like billing cycles, weekends, or holidays.
- Lack generalizability due to dataset inconsistency.

The dataset used in this study addresses these gaps by combining weather, time, and energy generation data, while also incorporating economic billing indicators. This enables the development of a forecasting model that is not only accurate and interpretable but also applicable in real-world, cost-sensitive energy management scenarios.

3 Methodology

The main intention of this study is to estimate hourly electricity prices based on energy production and consumption data along with weather and economic information. For this, the models selected were powerful deep learning models, chiefly Bidirectional LSTM networks augmented further through CNN layers. The broad methodology, which formed

the research methodology, roughly comprises a system incorporating data collection and cleansing, feature engineering, model training, and performance evaluation.

3.1 Data Sources

- Energy Dataset: 35064 hourly records from January 2015 up until January 2018. It has 29 features concerning energy generation which was been drop to 18 columns.
- Weather Dataset: Hourly weather data for Madrid, approximately 35K+ entries. This set includes meteorological variables which was drop to 10 relevant columns.

Python 3 was used for the entire analysis. Several supported libraries were applied at different stages of the project. Pandas and numpy were for data preprocessing. Then, matplotlib and seaborn came to aid the visualization part.

3.2 Preprocessing and Cleaning

Any time related columns in both datasets were first standardized into a consistent datetime format to ensure proper alignment and allow for time-based operations. Features having significantly large temporal discontinuities in their observations were dropped from the analysis (generation hydro pumped storage aggregated,etc.) because their irregular availability would have only brought noise into the forecasting.

After the preprocessing, the datasets were merged together on timestamps to establish a single unified time-aligned dataset of energy market indicators and relevant weather attributes. For some weather related variables, averaging was applied to reconcile the structural differences of meteorological data and energy datasets.(19)Following such harmonization, MinMax scaling was performed to classical numerical features such that values were rescaled to lie in the normalized range between 0 and 1.

While exploratory experiments considered PCA for dimensionality reduction, it was not much included in the initial setting but included in the hybrid deep learning approach could extract and refine relevant feature .In the developed architectures, CNN was combined with BiLSTM so that advantages of both spatial and time patterns could have been harnessed.

- Conv1D was used to detect sudden spikes and downward steps with the very short-term time series pattern in our data.
- MaxPooling1D reduced complexity while helping prevent overfitting.
- Bidirectional LSTM layers enabled learning from sequences both past and future in the data.
- Dense layers generated the output, i.e., predicted price.

Sequences of time steps and corresponding feature values:

1. Loss function was MSE.
2. Optimizer was Adam, with an automated learning rate adjustment by ReduceLROnPlateau.
3. Callbacks:
 - EarlyStopping, to avoid overfitting
 - ReduceLROnPlateau, to adjust learning rate during training

The dataset was divided into training and test sets using an 80 percent to 20 percent ratio . The model had a chance to be trained for up to 100 epochs and a batch size of 16 value.

Energy demand has been predicted using LSTM and GRU networks as complicated relationships between historical energy consumption trends and influencing factors such as weather, environment, and others are modeled. Because the two models can retain the historical data over a long period and process a given sequence, the models are extremely applicable when temporal dependence makes the prediction, thus forecasting in energy markets (16)

4 Design Specification

This design guideline outlines the conceptual architecture, operational methodology, and computational models deployed in the proposed hybrid deep learning scheme for electricity price forecasts. The section gives a quantitative and structural description of the system, i.e., functional requirements, basic architectural building blocks, and framework of operation used to fuse temporal and exogenous variables in market predictions.

4.1 System Design Overview

So-called constructional considerations are targeted at integrating different deep-learning paradigms, combining feature extraction through convolution and sequence modeling for short-term variation and long-term dependency modeling in electricity market data. The proposed architecture acts as one forecasting engine-large, ingesting energy generation data, consumption data, and weather data for day-ahead and real-time price predictions.

4.2 Main Architectural Components

The system consists of three primary computational modules:

The Data Preprocessing Subsystem loads, cleans, and synchronizes multi-source front-end datasets (energy generation data, weather information, historical prices). Normalization includes feature engineering based on time (hour, day of week, month) and the handling of missing/duplicated records. It also performs dimensionality reduction using PCA whenever required, with the goal of speeding up computation.

The Sequence Learning Subsystem, on the other hand, is mostly focused on learning temporal dependencies in the data; a handful of recurrent architectures are put in place:

1. **LSTM Network:** stacked long short-term memory layers capture long-term dependencies; dropout regularization is used.
2. **GRU Network:** is the computationally efficient variant with almost the same capacity to learn temporal sequences.
3. **CNN–BiLSTM Hybrid:** The convolutional layers extract short-term patterns that get input into a bidirectional LSTM for context-aware forecasting.

Evaluation Visualization Subsystem: Involves computing the performance metrics, which include RMSE and MAE, creating comparative visualizations of model performance.

4.3 Input Data and Performance Evaluation

The forecasting pipeline ingests synchronized hourly records consisting of:

- Energy Sector Data: Power generation by source, total load, and storage consumption.
- Weather Data: Hourly temperature, humidity, wind speed and direction, rainfall, and cloud cover.
- Prices: Past day-ahead and actual electricity prices.

Performance Evaluation: The primary measures of model performance are RMSE and MAE, which measure the average or mean-wise maximum deviation of the predicted prices from the true values in the corresponding scale.

4.4 System Framework

The entire system follows a pipeline architecture:

1. Data ingestion and cleaning.
2. Feature scaling and optional PCA transformation.
3. Sequence creation with configurable look-back windows.
4. Model training and hyperparameter tuning of architectures considered.
5. Forecast generation and evaluation.

Model implementations rely upon the Python language, with TensorFlow/Keras, NumPy, Pandas, and Scikit-learn. For data visualization, Matplotlib and Seaborn are used. For model training, it is, however, recommended to have the GPU environment, that is, with an NVIDIA Tesla T4 GPU or equivalent.

4.5 Algorithm Design

LSTM Algorithm: It handles sequences of data, remembering past states, while some gates provide mechanisms to let information flow into the system or vanish if unwanted, thus addressing the vanishing gradient problem.

GRU Algorithm: It works much as LSTM but is computationally faster due to fewer gates and therefore usually does not have a significant drawback on predictive capacity.

CNN–BiLSTM algorithm: The key step is that Conv1D temporal layers are responsible for extracting local temporals patterns. These are then passed onto LSTM layers that go bidirectionally between the two time directions, forward and backward.

4.6 Design Constraints

Key limitations are:

1. The requirements of huge historical data to train the high-capacity model.
2. Overfitting that may come into play anywhere in the model training if some method such as dropout or regularization is not used at some point appropriately along the way.
3. Demand for heavy computational resources that require GPU acceleration.
4. Sensitive towards feature scaling, and so each input should go through a preprocessing pipeline that is consistent throughout.

5 Implementation

The final implementation phase focused on integrating meteorological data and energy consumption data to construct an enriched dataset to serve the purposes of modeling and analyzing in greater depth.

Implemented in the case under study by relying on a plethora of open source data science ecosystems such as pandas, NumPy, and Jupyter Notebook, the procedure first integrated two main sources gathered from either meteorological measurements (weather features.csv) or energy generation and pricing (energy dataset.csv). Once the timestamps were standardized in terms of time zones, the merging was done upon those synchronized time-based keys to assure a perfect alignment of the two datasets. The above described manipulation produced a single powerful and well structured database having 36 columns (after including new features) and containing different types of information, including typical weather variables such as temperature, relative humidity, and wind speed; records of energy output by source; and derived time features such as time of the day, day of the week, month, and quarter.

Comprehensive preprocessing was done to get the data ready for modeling:

1. To improve the time-series context, new temporal features were included.
2. To make them compatible with numerical operations, data types were standardized, particularly by changing pertinent columns to float type.
3. In order to maintain data integrity, missing values were handled appropriately.
4. The use of feature engineering was essential. To facilitate temporal modeling, temporal information was taken from the timestamp. New measures were created by calculating weekend indicators and monthly electricity billing aggregates.

In case of CNN-BiLSTM, prior to performing PCA, MinMaxScaler acted to standardize the dataset; this operated to rescale all intervening features to a common scale from 0 to 1. This helped prevent any one variable from dominating due to their differing scales and helped in stabilizing and speeding up the convergence of models during training. Visualization techniques were at the very heart of the process to enable understanding the relation, seasonal trends, and interactions between various energy market variables and weather conditions. By visualizing these patterns specifically, demand rising due to extreme temperatures, or price peaks that reoccur by season, we were able to develop intuition on the features to include and on the architecture of the models.

We tested deep learning methods with other conventional sequential learning methods to evaluate their strengths. The most sophisticated architecture created was the CNN-BiLSTM hybrid architecture, which first processed the data through convolutional layers to learn local feature patterns and then through bidirectional long short-term memory layers to learn time dependencies, both forward and backward. To further improve learning and prevent overfitting, the following were added: MaxPooling1D layers, Dropout layers, and fully connected Dense layers. This combination enabled the model to more effectively learn from the spatial and temporal cues, making it a good fit for the complicated dynamics of electricity price forecasting.

In parallel to the hybrid deep learning approach, the dataset was also put to use to train more conventional RNNs, namely LSTMs and GRUs. These were considered the baseline to judge how much convolutional layers and the bidirectional mechanism added to the hybrid CNN-BiLSTM architecture. By comparing the performance results, one reached a direct assessment of how much better the advanced design was compared to that of more traditional sequence-based models.

The deep learning models were given especially careful treatment, with focus placed on training strategies that allowed efficiency and reliability. EarlyStopping, in particular, was used to watch the validation loss and stop training when no further improvement was seen, thereby avoiding unnecessary computation and overfitting. Further, ReduceLROnPlateau was engaged to automatically reduce the learning rate when progress slowed, allowing the models to finetune their learning and converge more effectively.

Having cleaned and well-structured datasets at this stage, combined with a range of trained models whose performance was rigorously tested against suitable regression benchmarks, meant not only a set of precise forecasting tools were in hand but a solid foundation upon which to build future fine-tuning and experimentation.

6 Evaluation

From the plotted graphs in fig 6.1, a clear inverse relationship emerged between wind speed and electricity prices: as wind speed increases, electricity costs tend to decrease. This pattern is likely due to the higher contribution of low cost wind energy to the power grid during periods of strong winds. When more wind power is available, reliance on more expensive generation sources—such as fossil fuels or peak demand gas plants—is reduced, leading to a noticeable drop in overall electricity prices.




Fig 6.1. Wind vs Energy Price

As shown in fig 6.2, the hourly price trends revealed a distinct daily rhythm. Electricity prices are at their lowest during the early morning hours, particularly between 2:00 and 5:00 AM, when overall demand is minimal. Prices then rise sharply starting around 6:00 AM, peaking between 9:00 and 10:00 AM as businesses open and residential energy use increases. A second, often higher, peak occurs in the evening between 7:00 and 8:00 PM, coinciding with the time most people return home, cook, and use household appliances. These fluctuations closely mirror typical daily demand cycles—morning activity driving the first spike, and post work home life driving the second.

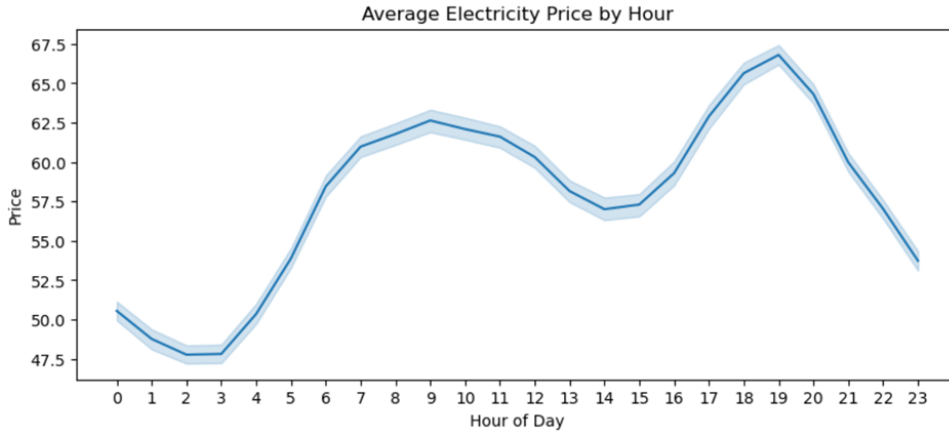


Figure 6.2. Hourly Average Price

As shown in fig 6.3, looking at the energy generation mix, the system remains heavily dependent on non-renewable sources. The largest contributors are nuclear power (22.2 percent), fossil gas (20.0 percent), fossil hard coal (15.1 percent), and wind onshore (19.4 percent). While renewable sources such as wind, solar, hydro, and biomass are present, their contributions are generally smaller and more variable. Overall, non-renewable account for 61.1 percent of total energy generation, compared to 38.9 percent for renewables. This dominance of fossil fuels and nuclear power highlights the sector’s ongoing reliance on traditional energy sources.

Although the share of renewable energy has been gradually increasing, the current energy mix still leaves electricity prices vulnerable to fossil fuel market volatility and the supply chain risks associated with these resources. Until renewable capacity expands further and becomes more consistent, these dependencies will continue to influence both price stability and long-term energy security.

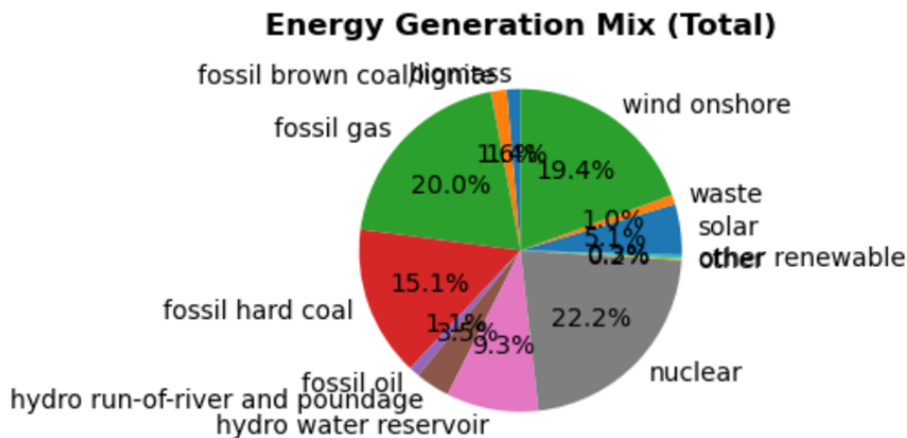


Figure 6.3. Energy Sources

6.1 Case Study 1: LSTM-Based Forecasting

The first experiment involved analyzing a blend of historical energy generation data, with detailed weather records thrown in, all to properly ascertain the prices for future energy markets (usually termed as price actual). The actual forecasting was carried out using the LSTM model, a special type of neural network that codes for time series well due to

its long-term dependency capture ability: letting the model forecast the price of the next hour on past patterns in the preceding 24-hour sequence of data.

Before making an array out of all data into the model, there were a couple of preprocessing steps to guarantee peak performance: from setting the sequential input format required for time series modeling to scaling all features of the input data using MinMax scaling, which improves learning efficiency of the model. The LSTM architecture was made up of two layers stacked on top of each other, both containing LSTM units of 50 in number, in order to help the model capture more complex temporal features. This was followed by a dense layer with leucine activation to introduce non-linearity and dropout layers introduced to prevent overfitting. The model was compiled with the adam optimizer.

After training and evaluation, the model showed promising performance. The training and validation loss curves converged well, implying that meaningful patterns were learned from the data. Quantitatively, the MAE was 4.47, and 5.99 was the RMSE. Visual comparisons of the predicted price curve and the true price curve were carried out; in general, the model was good at keeping track of price trends while showing some deviation at certain points. Slight overfitting, however, did start to develop in training, indicating that the model’s generalization ability might be improved with further regularization or simplified network complexity.

In the end, the LSTM model was a strong and reliable baseline for time series prediction for the energy price case, providing useful insight in conjunction with a future potential for refinement.




Figure 6.4. LSTM Model Results

6.2 Case Study 2: GRU-Based Forecasting

The second experiment considered the use of gated recurrent units, an alternative type of recurrent neural network architecture that has been attracting attention in recent years. GRUs are quite similar to LSTM networks, actually, but are mostly faster to train and require fewer parameters, which are very suitable when computational resources are limited or faster iteration cycles are needed; that is to say, a simplification of the internal structure of gates allows them to encode important temporal dependencies in time series with less computation complexity compared to LSTMs.

For direct comparison with the previous experiment, the GRU model was built in an architecture very close to the LSTM one, i.e., two stacked recurrent layers with 50 units in each layer, followed by a dense layer whose function is to transform the learned features into predictions, and dropout layers to help prevent overfitting. All preprocessing steps were kept the same — sequence creation, MinMax scaling — training tactics, callbacks, and optimization settings — were exactly the ones used in the LSTM experiment. Ensuring this consistency gave greater confidence in attributing any difference in performance to

architecture rather than some other experimental variable.

Results from the GRU model were encouraging. It obtained a Root Mean Square Error (RMSE) of 5.75 and a Mean Absolute Error (MAE) of 4.27 — figures quite close to those of the LSTM. Besides, the convergence was somewhat faster, meaning that GRU required fewer training epochs to reach optimal performance. A visual check on the prediction plots showed that the GRU outputs were in strong agreement with the observed ones, accurately tracking short-term movements and long-term price trends.

Considering that these were moderate improvements over LSTM, the savings in computation afforded by GRUs make them appealing choices, especially when forecasting performance in the real world is resource constrained and time constrained — e.g., performing forecasts on low power devices or generating updates fast under very dynamic market conditions. In short, the GRU showed that it could deliver near identical predictive accuracy to the LSTM while demands less time and processing power, making it a practical alternative for time series forecasting in energy price model.




Figure 6.5. GRU Model Results

6.3 Case Study 3: CNN-BiLSTM Hybrid Model

The third case study used a hybrid deep learning architecture, combining the power of CNNs and BiLSTM technology. The rationale behind this design was to have a model that could simultaneously grasp localized short-term representations in the data and long-range sequential dependencies; in other words, the capacity of CNNs to lump together local motifs with the temporal view of recurrent networks.

Initially, the method applied a 1D CNN with a kernel size of 3, which allowed it to examine the input time series in smaller overlapping windows. This step was rather effective in extracting short-term changes and local trends, for instance, those short-lived spikes or dips seen in energy prices due to a rapid change in supply or demand. Post this convolution operation, MaxPooling was applied to down sample and throw away irrelevant information to a certain extent so that the network would largely concentrate on major features and, interestingly, reduce computational cost. The condensed feature maps were passed to a bidirectional LSTM layer. While a unidirectional LSTM deals with information flow in a single temporal direction, its bidirectional counterpart processes the sequence both forward and backward. With regard to highly dynamic markets like energy pricing, this dual perspective allows the model to consider how past and future contexts could influence each point in time. Then dropout regularization was applied so as to reduce overfitting and improve generalization, followed by the dense layers that translated the learned features into price predictions.

When training and evaluation were performed, the hybrid models outshone those provided by LSTM or GRU networks used on their own. It gave a root mean square error (RMSE)

of 5.18 and mean absolute error (MAE) of 4.10, which were the lowest achieved by all tested architectures. Looking at time-series plots, it is evident that this model retained greater stability in its predictions, showing lesser random deviations than previous models. Moreover, the scatter plots comparing predicted and actual values placed points closer in the ideal diagonal line, which signifies stronger coherence between the forecast and the reality.

In all, the experiment set forth an example of the capacity of combining convolutional and recurrent layers for time series forecasting in volatile situations such as the energy markets, which contain short-living random fluctuations and long-term dependencies. The hybrid model, capturing local patterns through CNNs and enhancing them using bidirectional temporal reasoning by LSTMs, proved to be highly flexible and robust, thus setting the bar high for any further experimentation in the area.




Figure 6.6. CNN-BiLstm Model Results

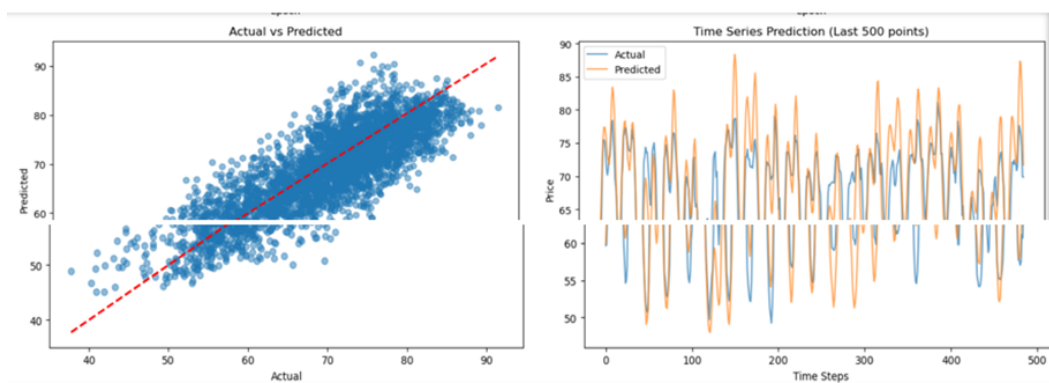


Figure 6.7. Scatter Plot And Time Series Plot

6.4 Experiment / Case Study 4: PCA and Seasonal Decomposition

The fourth experiment followed a slightly different path, employing statistical analysis techniques to complement and perhaps even support deep learning models. These included seasonal decomposition and PCA to distill insights into the very structure of the data and thereby enhance the preprocessing steps.

PCA analysis was applied to the concatenation of weather and energy datasets in order to perform dimensional reduction: converting many correlated variables into a handful of principal components. Essentially, the goal was to construct a lower-dimensional representation that would retain 95 percent of the variance in the original data, so that the majority of potentially relevant information was never compromised. By excluding

the noise and redundancy present in the data, the step offered a good pre-processing mechanism since it would assist the models with processing cleaner data and save time during the training phase. Furthermore, this step helped in better focusing the models on the actual underlying patterns, thereby improving performance and ease of interpretability. Simultaneously, the additive seasonal decomposition model was also fit to the hourly aggregated prices actual time series. This technique decomposes a time series into three components: trend, seasonal, and residuals (or noise). It is in this way that the decomposition showed clear, expected monthly seasonal patterns in energy prices: the peaks and valleys almost perfectly coincided with known cycles in energy demand, e.g., less consumption in months of temperate weather, more in extreme months of cold or heat. Such seasonal behavior was a reflection of the regular flips in market dynamics due to changes in weather or cyclical economic activities.

This kind of statistical experiment went beyond simply confirming the preexisting patterns and actually gave us a data-oriented footing for helping to improve future models toward an explicit seasonality leaning. Such a constructive way of learning about the temporal structure of the given data naturally paves the way for placing this information into the forecast models that would then be able to anchor upon these recurring trends, giving us more precise and robust forecasting capabilities. This experiment's emphasis was on the combination of the traditional statistical approach and modern machine learning in order to draw more interesting insights and stronger predictive models.

6.5 Discussion

The comparative study offers new and valuable insights into the performance of the different deep learning architectures when predicting the price of energy. One key takeaway is that the GRU model attains more accuracy comparable to that of the more complex LSTM but can be trained much faster and with fewer resources, which makes it suitable for real-time applications or computing power challenged environments.

On that side, a hybrid CNN-BiLSTM proved the power of an integration of approaches that dealt with spatial modeling and temporal modeling. While convolutional layers are good at extracting short-term, local patterns, the bidirectional LSTMs capture longer-range dependencies in both directions of time. Such a hybrid architecture outperformed the LSTM and GRU standalone models in terms of accuracy and resilience, meaning its predictions were more reliable even when the energy markets were highly volatile.

From classical statistics' point of view, the applications of seasonal decomposition and Principal Component Analysis (PCA) confirmed that the underlying information contains predictable structures, such as seasonal trends and correlations between variables. Such a cardinal insight allows a statement to be made regarding how inconstantly hybrid methods will utilize explicitly the seasonal components, thereby reaching further in accuracy and robustness in forecasting. The positive results came with some set backs during the research. An example of this is the fact that all models were trained with alphaing parameters fixed; as such, some of the models may not have shown their full capacity. The geographic scope was limited also as only weather data from Madrid was taken into consideration, limiting the applicability of the findings to other areas with differing climatic conditions or market situations. Another highly relevant aspect for adoption has to do with the (unfortunately neglected) explainability of some models: in simple words, it is difficult to build trust and make a decision for energy systems when the basis for some decisions leaned upon some complex mathematical algorithms that

Models	RMSE	MEA
LSTM	5.99	4.47
GRU	5.75	4.27
CNN-BiLstm	5.18	4.10

Figure 2: Model Results

cannot be explained to a market operator.

Placing ideas forward, several promising research directions could be opened. First, one could systematically employ hyperparameter optimization to further tune model architectures to improve their performance. Then, incorporating external factors such as holidays, policy changes, or economic indicators could provide a richer context and thereby strengthen the models’ ability to reflect more complex market behavior. More recent architectures like attention mechanisms or transformer-based models can be incorporated to enhance prediction robustness and adaptability. Finally, enhancing model explainability through explainability tools like SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model agnostic Explanations) would enhance the interpretability and reliability of such models for energy stakeholders.

All things aside, this paper points to the area of improvement and the enormous potential of deep learning to enhance time-series forecasting in energy markets. It paves the way for possible future breakthroughs that can lead to more efficient, accurate, and interpretable prediction models that are tailored to the dynamic and complex nature of energy systems.

7 Conclusion and Future Work

The purpose of this study was to determine whether a hybrid deep learning model that combines Convolutional Neural Networks (CNN) and Bidirectional Long Short-Term Memory (BiLSTM) networks could estimate power prices more accurately than conventional sequence-based models like LSTM and GRU. The main goal was to show that the hybrid design is more effective in capturing long-range relationships as well as local temporal patterns in complicated time-series data that are produced from weather, energy generation, and consumption.

This was accomplished by combining energy mix data (containing data from different generation sources including solar, wind, fossil fuels, and nuclear), electrical load, meteorological metrics (such temperature, wind speed, humidity, and pressure), and historical power pricing to create a comprehensive dataset. In order to capture seasonal and cyclical effects, features including the month, day of the week, and hour of the day were designed. Three deep learning models—a hybrid CNN-BiLSTM model, a GRU network, and a regular LSTM model—were trained and assessed. Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) were used to gauge performance.

The results shown in Fig 2 unequivocally demonstrate that the CNN-BiLSTM hybrid model performed better than the LSTM and GRU models on every evaluation metric. The hybrid model demonstrated greater prediction accuracy and generalization capabilities by achieving the lowest RMSE and MAE as well as the greatest R2 score.

Despite its success, the study has limitations. The model’s accuracy is highly dependent on the quality and granularity of the input data. It does not currently account for external shocks such as policy changes, etc. Moreover, while the model achieves high performance, it lacks interpretability.

Using the CNN-BiLSTM model in real-time settings and possibly integrating it with edge devices or smart grid infrastructure for localized forecasting is a significant approach. Future research should also use model explainability tools, like SHAP or LIME, to determine which elements have the most influence on predictions in order to encourage wider use. Additionally, this model's commercial viability could be improved by attempting to apply domain adaptation or transfer learning to other geographical areas.

To sum up, this study has effectively shown that a CNN-BiLSTM hybrid model provides a notable improvement in electricity price forecasting compared to conventional deep learning models, paving the way for more reliable, scalable, and intelligent energy sector solutions.

References

- [1] Ballester, C. and Furió, D. (2024). Renewable generation and wholesale electricity pricing components in the iberian market: A machine learning approach, *Green Finance* **6**(2): 201–223.
URL: <https://www.aimspress.com/aimspress-data/gf/2024/2/PDF/GF-06-02-014.pdf>
- [2] Belenguer, E., Segarra-Tamarit, J., Pérez, E. and Vidal-Albalade, R. (2025). Short-term electricity price forecasting through demand and renewable generation prediction, *Mathematics and Computers in Simulation* **229**: 350–361.
- [3] Cerjan, M., Krčelj, I., Vidak, M. and Delimar, M. (2013). A literature review with statistical analysis of electricity price forecasting methods, *2013 IEEE EUROCON*, pp. 756–763.
- [4] Chen, S., Fang, X. and Khayatnezhad, M. (2024). Forecasting for electricity demand utilizing enhanced inception-v4 using improved osprey optimization, *Scientific Reports*. Open access.
- [5] Ghosh, A. and Shah, A. (2024). Digitalization in the renewable energy sector: A literature review, *Energies* **17**(9): 1985.
- [6] Haas, R. (2022). On the economics of storage for electricity: Current state and future market design prospects, *WIREs Energy and Environment* **11**(3). Review covers relevance of storage options, costs, economics, welfare effects, and issues of electricity market design.
- [7] Jedrzejewski, D., Marcjasz, G. and Weron, R. (2022). Electricity price forecasting: Comprehensive review and implications for the energy transition, *Energy Economics* **112**: 106127.
- [8] Khan, A., Iqbal, N., Rizwan, A., Ahmad, R. and Kim, D. (2021). An ensemble energy consumption forecasting model based on spatial-temporal clustering analysis in residential buildings, *Energies* **14**(11): 3020.
- [9] Marcjasz, G., Narajewski, M., Weron, R. and Ziel, F. (2022). Distributional neural networks for electricity price forecasting, *Applied Energy* **325**: 119841.

- [10] Matsumoto, T. (2021). *One-week-ahead electricity price forecasting using weather forecasts, and its application to arbitrage in the forward market: an empirical study of the Japan Electric Power Exchange*, Doctoral dissertation, Kanazawa University.
- [11] Mosquera-López, S. (2024). The impact of weather variability on renewable energy consumption: Insights from explainable machine learning models, *Energy Economics* **129**: 107394.
- [12] Neumann, D., Riedl, T., Meier, P. and Hildebrandt, T. (2023). Using weather data in energy time series forecasting, *Energy Informatics* **6**(1): 10.
- [13] Pata, U. (2025). How to progress towards sustainable development by leveraging renewable energy sources, technological advances, and human capital, *Renewable Energy* p. 122367.
- [14] Ramnath, R., Thomas, B., Ajay, S. and Shukla, R. (2024). Predicting household electricity consumption using hybrid ensemble models with survey-based and historical features, *Sustainability* **16**(13): 5417.
URL: <https://www.mdpi.com/2071-1050/16/13/5417>
- [15] Tanaka, R. (2022). The impact of weather changes on the supply and demand of electricity in germany: A structural equation modeling approach, *Energy Economics* **109**: 105958.
- [16] Unsal Çelimli, D., Aksöz, A., Oyucu, S., Guerrero, J. and Guler, M. (2024). A comparative study of ai methods on renewable energy prediction for smart grids: Case of turkey, *Sustainability (Switzerland)* **16**(7).
- [17] Vengatesan, S., Jayakumar, A. and Sadasivuni, K. K. (2024). Fcev vs. bev—a short overview on identifying the key contributors to affordable & clean energy (sdg-7), *Energy Strategy Reviews* **53**: 101380.
- [18] Wikipedia contributors (2025). Grid energy storage, https://en.wikipedia.org/wiki/Grid_energy_storage. Accessed: 2025-08-10.
- [19] Yu, B., Li, J., Liu, C. and Sun, B. (2022). A novel short-term electrical load forecasting framework with intelligent feature engineering, *Applied Energy* **327**: 120089.
- [20] Zachmann, G. (2023). The design of the european electricity market, *Study*, European Parliament, Policy Department for Economic, Scientific and Quality of Life Policies, Directorate-General for Internal Policies.