

Real-Time Electricity Price Prediction and Trading Signal Generation Using Ensemble Tree-Based Machine Learning Models: A Comparative Study on the Spanish Electricity Market

Yirui Liu*

College of Computer Science and Technology, Shandong University of Technology, Zibo 255049, Shandong, China

*Corresponding author: Yirui Liu, 2110621759@qq.com

Copyright: © 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0), permitting distribution and reproduction in any medium, provided the original work is cited.

Abstract: Accurate real-time electricity price forecasting is of critical importance for market participants seeking to optimize energy procurement, dispatch scheduling, and arbitrage strategies in liberalized electricity markets. However, existing forecasting approaches suffer from several key limitations: (1) conventional statistical models fail to capture the complex nonlinear interactions among generation mix, load demand, and temporal variables that collectively drive price dynamics; (2) single-model approaches lack robustness and are sensitive to overfitting, limiting their generalizability across diverse market conditions; (3) the interpretability of black-box prediction models remains insufficient, hindering the practical deployment of data-driven forecasting systems in operational decision-making. To address these challenges, this study proposes a comprehensive machine learning framework based on six tree-based ensemble models for hourly electricity price prediction in the Spanish electricity market. The proposed framework introduces three key contributions: (1) a systematic feature engineering pipeline incorporating lagged price variables, rolling statistics, and calendar-based temporal encodings; (2) a rigorous comparative evaluation of Decision Tree, Random Forest, Extra Trees, Gradient Boosting, XGBoost, and LightGBM under identical experimental conditions; (3) a SHAP-based interpretability analysis that quantifies feature contributions and interaction effects at both global and local levels. Experimental results on the ENTSO-E Spanish market dataset demonstrate that XGBoost achieves the best overall predictive performance, with an R^2 of 0.9660 and MAE of 1.5631 €/MWh.

Keywords: Electricity price forecasting; Ensemble learning; Gradient boosting; XGBoost; LightGBM; SHAP interpretability; Spanish electricity market

Online publication: March 26, 2026

1. Introduction

In recent years, electricity price forecasting (EPF) has emerged as a fundamental problem in power system management and energy market optimization ^[1]. Accurate real-time price prediction enables market participants to develop effective bidding strategies and maximize economic returns. With the rapid integration of renewable energy, price dynamics have become increasingly volatile and nonlinear, making reliable forecasting more challenging than ever.

Although significant progress has been made using statistical and machine learning approaches, existing methods still suffer from limited generalizability, insufficient model interpretability, and poor performance under high-volatility conditions ^[2-6]. Single-model frameworks, in particular, fail to capture the complex interactions among generation mix, load demand, and temporal variables ^[7,8].

To address these challenges, this paper proposes a comprehensive tree-based ensemble learning framework for hourly electricity price forecasting on the Spanish ENTSO-E market dataset, further incorporating SHAP-based interpretability analysis to reveal key price-driving features ^[9-11].

The main contributions are as follows:

- (1) A systematic comparison of six tree-based models under identical experimental conditions;
- (2) A feature engineering pipeline integrating lagged price variables, rolling statistics, and temporal encodings;
- (3) A SHAP-based interpretability analysis quantifying global feature contributions and pairwise interaction effects.

2. Experimental results and analysis

This chapter presents a comprehensive evaluation of the six tree-based machine learning models on the Spanish hourly electricity price dataset. The experimental analysis is organized into three components: prediction curve comparison to assess temporal tracking fidelity, quantitative performance metric evaluation to rank model accuracy, and a visualization of trading signals generated by the best-performing model. All experiments are conducted under a unified protocol with an 80/20 chronological train-test split and evaluated using MAE, RMSE, R^2 , and MAPE.

2.1. Prediction curve comparison

Figure 1 illustrates the temporal alignment between the actual electricity prices and the predicted values generated by each of the six models over 500 consecutive test time steps. Across all subplots, the predicted curves closely track the black actual price curve, faithfully reproducing both the short-term oscillations and the broader downward trend observed between time steps 300 and 420. The Decision Tree exhibits the largest deviations, particularly at price spikes, consistent with its lower R^2 of 0.9362. In contrast, the four ensemble models, Random Forest, Extra Trees, Gradient Boosting, XGBoost, and LightGBM, achieve substantially tighter fits, with their predicted curves nearly overlapping the actual series throughout the evaluation window.

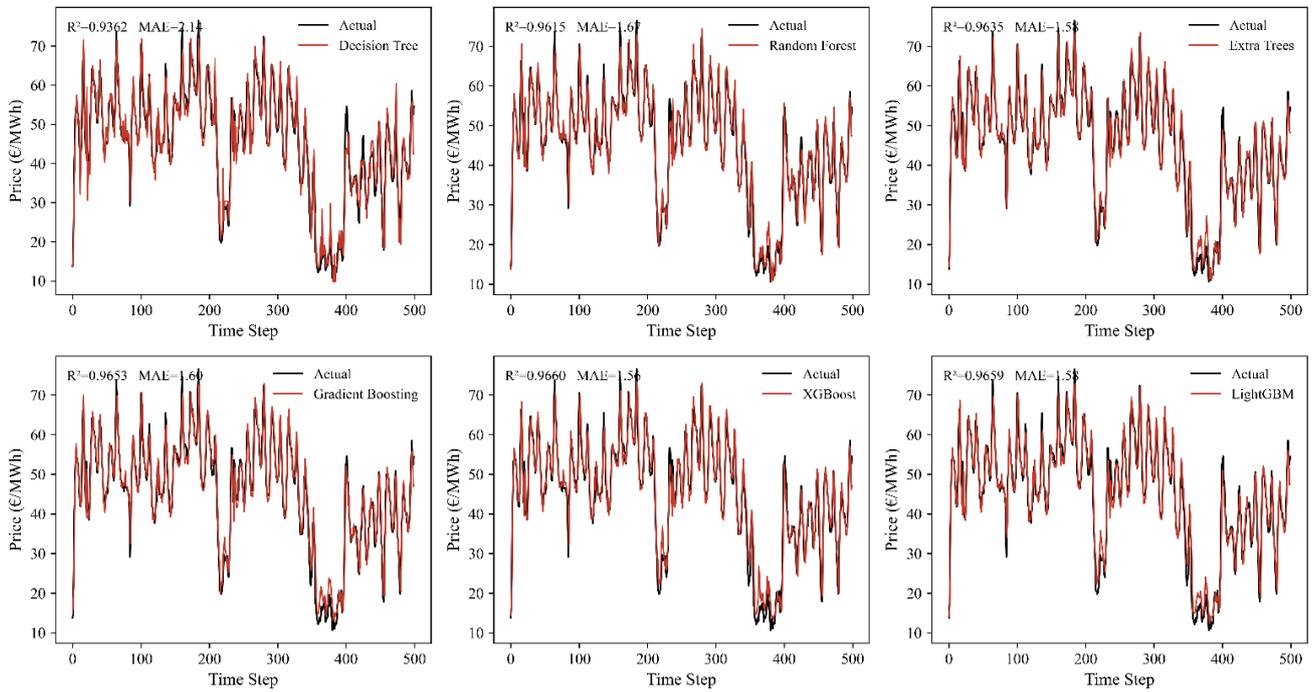


Figure 1. Comparison of actual and predicted electricity price curves for six tree-based machine learning models over 500 consecutive test time steps.

Each subplot displays the black actual price series against the red predicted series, with the corresponding R^2 and MAE values annotated in the upper left corner. The results demonstrate that ensemble models, particularly XGBoost and LightGBM, achieve substantially closer alignment with the actual price trajectory compared to the single Decision Tree baseline, effectively capturing both short-term volatility and longer-term price trends in the Spanish electricity market.

2.2. Quantitative performance comparison

Table 1 reports the MAE, RMSE, R^2 , and MAPE values achieved by each of the six models on the held-out test set. A clear and consistent performance hierarchy emerges across all four metrics: the single Decision Tree yields the weakest results ($R^2 = 0.9362$, MAE = 2.1368 €/MWh), while all five ensemble models substantially surpass it. Among the ensemble methods, XGBoost achieves the best overall performance with $R^2 = 0.9660$ and MAE = 1.5631 €/MWh, marginally outperforming LightGBM ($R^2 = 0.9659$) and Gradient Boosting ($R^2 = 0.9653$). Extra Trees and Random Forest deliver competitive results, confirming that Bagging-based methods also provide strong predictive accuracy. The narrow performance gap among the four boosting and bagging ensembles suggests that the primary determinant of accuracy in this task is the adoption of an ensemble strategy rather than the specific boosting algorithm employed.

Table 1 was evaluated using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Coefficient of Determination (R^2), and Mean Absolute Percentage Error (MAPE). Lower MAE, RMSE, and MAPE values indicate higher prediction accuracy, while a higher R^2 value reflects stronger explanatory power. XGBoost achieves the best performance across all metrics.

Table 1. Quantitative performance comparison of six tree-based machine learning models on the Spanish electricity price test set

Model	MAE	RMSE	R ²	MAPE (%)
Decision Tree	2.1368	2.9128	0.9362	3.5430
Random Forest	1.6721	2.2632	0.9615	2.7878
Extra Trees	1.5810	2.2036	0.9635	2.6662
Gradient Boosting	1.5974	2.1489	0.9653	2.6913
XGBoost	1.5631	2.1249	0.9660	2.6579
LightGBM	1.5769	2.1289	0.9659	2.6582

2.3. Trading signal identification

Figure 2 presents the buy and sell signals generated by the best-performing XGBoost model over the 500-step test window, overlaid on both the actual and predicted price curves. Local price minima identified in the predicted series are marked as buy signals (upward-pointing triangles), while local maxima are designated as sell signals (downward-pointing triangles). The predicted curve closely shadows the actual price trajectory throughout the evaluation period, resulting in trading signals that are well-aligned with the true price turning points. The buy signals predominantly coincide with genuine price troughs, and the sell signals consistently appear at or near true price peaks, confirming the practical utility of the XGBoost predictions for intraday electricity arbitrage strategies. The high density of accurately positioned signals across the full test window demonstrates that the model captures both the fine-grained oscillatory structure and the broader price cycles of the Spanish electricity market with sufficient fidelity to support real-time trading decision support.

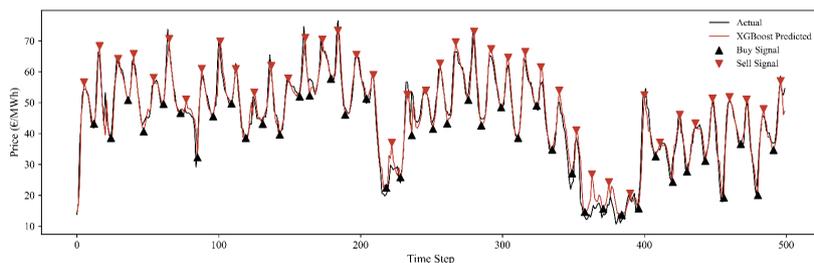


Figure 2. Trading Signal Identification.

3. SHAP-based interpretability analysis

Building upon the predictive performance, this section employs SHapley Additive exPlanations (SHAP) to systematically investigate the feature importance and interaction effects underlying the XGBoost model. SHAP decomposes each prediction into additive contributions from individual features, providing a theoretically grounded and model-agnostic framework for interpreting the learned decision logic.

3.1. SHAP dependence plots

Figure 3 presents SHAP dependence plots for the six most influential features identified in the XGBoost model. Each subplot displays the relationship between a feature’s raw value on the horizontal axis and its corresponding SHAP contribution on the vertical axis, with point color encoding a secondary interaction feature. Panel (a)

reveals a near-linear positive relationship between price day ahead and its SHAP value, confirming that the day-ahead price is by far the dominant predictor of actual electricity prices, contributing SHAP values ranging from approximately -30 to $+40$ €/MWh. Panel (b) shows that generation from fossil hard coal exerts a moderately negative influence at intermediate generation levels, reflecting the merit-order effect in which coal plants occupy a mid-stack position. Panel (c) indicates that nuclear generation contributes small but consistent negative SHAP values across its observed range, consistent with its role as a low-cost baseload source that suppresses spot prices. Panels (d) through (f) demonstrate that hydro run-of-river output, total load forecast, and total load actual all exhibit positive SHAP contributions at higher values, reflecting demand-driven price escalation under high-load conditions.

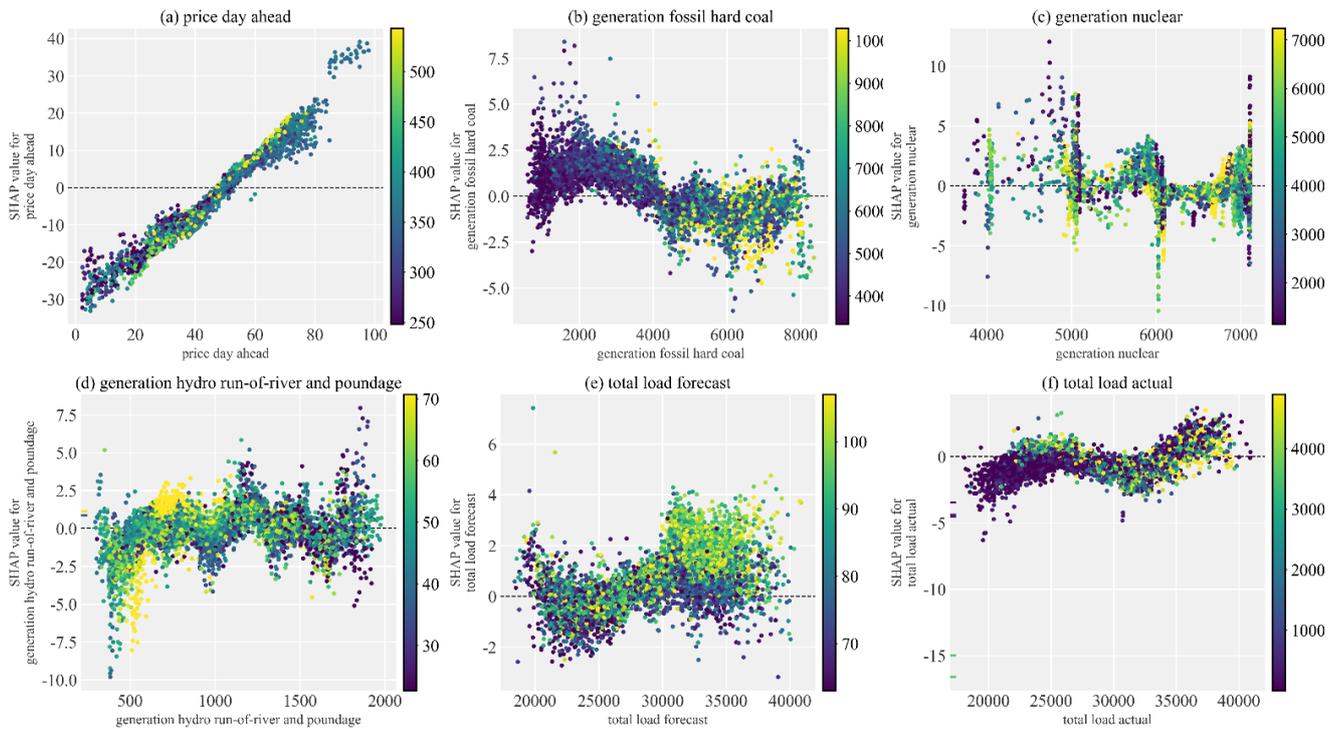


Figure 3. SHAP dependence plots for the six most influential input features of the XGBoost electricity price prediction model, including price day ahead, generation fossil hard coal, generation nuclear, generation hydro run-of-river and poundage, total load forecast, and total load actual.

Each point represents one test sample, with the horizontal axis showing the feature value and the vertical axis showing the corresponding SHAP contribution to the predicted price. Point colors encode a secondary interacting feature value. The plots reveal the direction, magnitude, and nonlinearity of each feature’s marginal effect on electricity price predictions.

3.2. SHAP interaction dependence plots

Figure 4 presents SHAP interaction dependence plots for three key feature pairs, quantifying the conditional interaction effects that cannot be captured by marginal SHAP values alone. Panel (1) reveals a pronounced nonlinear interaction between total load forecast and total load actual: at low forecast levels the interaction value is negative, transitioning to strongly positive beyond approximately 30,000 MW, indicating that the price-elevating

effect of high load is amplified when both forecasted and realized demand are simultaneously elevated. Panel (2) shows that the interaction between generation biomass and price day ahead is relatively weak and dispersed across the mid-range of biomass output. Panel (3) demonstrates that generation nuclear interacts with price day ahead in a concentrated manner near the 4,500–5,000 MW nuclear output band, where interaction values cluster tightly around zero, suggesting that nuclear generation’s price-suppressing effect is partially moderated by the prevailing day-ahead price level.

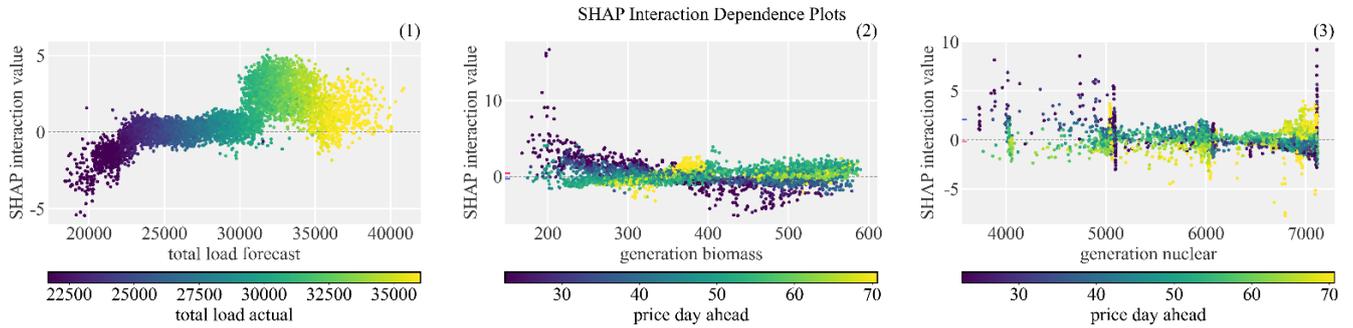


Figure 4. SHAP interaction dependence plots for three selected feature pairs: (1) total load forecast interacting with total load actual, (2) generation biomass interacting with price day ahead, and (3) generation nuclear interacting with price day ahead. Vertical axes show SHAP interaction values and point colors encode the secondary feature.

4. Conclusion

This study proposed a comprehensive tree-based machine learning framework for real-time electricity price forecasting and trading signal generation in the Spanish electricity market. Six models, Decision Tree, Random Forest, Extra Trees, Gradient Boosting, XGBoost, and LightGBM, were systematically constructed, trained, and evaluated under identical experimental conditions. Experimental results demonstrated that all ensemble models substantially outperformed the single Decision Tree baseline, with XGBoost achieving the best overall performance across all four evaluation metrics ($R^2 = 0.9660$, MAE = 1.5631 €/MWh, RMSE = 2.1249, MAPE = 2.6579%). SHAP-based interpretability analysis further revealed that price day ahead is the dominant price driver, while generation mix and load variables contribute nonlinear secondary effects that interact in economically meaningful ways. The proposed framework also demonstrated practical utility in generating actionable buy and sell signals aligned with actual price turning points. Future work will explore the integration of weather features and cross-market data to further improve forecasting accuracy, the incorporation of deep learning architectures such as Transformers for long-horizon price prediction, and the development of a real-time deployment pipeline for operational electricity trading support systems.

Disclosure statement

The author declares no conflict of interest.

References

- [1] Jędrzejewski A, Lago J, Marcjasz G, et al., 2022, Electricity Price Forecasting: The Dawn of Machine Learning.

IEEE Power and Energy Magazine, 20(3): 24–31.

- [2] Wang P, Xu K, Ding Z, et al., 2022, An Online Electricity Market Price Forecasting Method Via Random Forest. IEEE Transactions on Industry Applications, 58(6): 7013–7021.
- [3] Shah R, Shah H, Bhim S, et al., 2021, “Short-Term Electricity Price Forecasting using Ensemble Machine Learning Technique. 2021 1st International Conference in Information and Computing Research (iCORE), Manila, Philippines, 145–150.
- [4] Sahoo S, Swain S, Dash R, 2023, An Analysis of Machine Learning Methods for Electricity Price Forecasting. 2023 1st International Conference on Circuits, Power and Intelligent Systems (CCPIS), Bhubaneswar, India, 1–5.
- [5] Alkawaz A, Abdellatif A, Kanesan J, et al., 2022, Day-Ahead Electricity Price Forecasting Based on Hybrid Regression Model. IEEE Access, 2022(10): 108021–108033.
- [6] Mubarak H, Ahmad S, Hossain A, et al., 2023, Short-Term Electricity Price Forecasting Using Interpretable Hybrid Machine Learning Models. 2023 IEEE IAS Global Conference on Renewable Energy and Hydrogen Technologies (GlobConHT), Male, Maldives, 1–6.
- [7] Arya K, Vijaya Chandrakala K, 2021, Machine Learning Based Prediction and Forecasting of Electricity Price During COVID-19, 2021 IEEE International Power and Renewable Energy Conference (IPRECON), Kollam, India, 1–6.
- [8] Han L, Ban C, Zhang C, et al., 2024, Electricity Price Forecasting in Power Markets Based on Machine Learning, 2024 IEEE 4th International Conference on Digital Twins and Parallel Intelligence (DTPI), Wuhan, China, 91–96.
- [9] Heijden T, Lago J, Palensky P, et al., 2021, Electricity Price Forecasting in European Day Ahead Markets: A Greedy Consideration of Market Integration. IEEE Access, 2021(9): 119954–119966.
- [10] Yorat E, Zor K, Özbek N, 2023, Day-Ahead Electricity Price Forecasting Using Artificial Intelligence-Based Algorithms, 2023 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT), Sakheer, Bahrain, 121–126.
- [11] Yildirim B, Yildiz S, Turkoglu A, et al., 2023, Evaluating LMP Forecasting with LSTM Networks: A Deep Learning Approach to Analyzing Electricity Prices During Unpredictable Events, 2023 5th Global Power, Energy and Communication Conference (GPECOM), Nevsehir, Turkiye, 477–482.

Publisher’s note

Bio-Byword Scientific Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.