

A Few-Shot LLM Framework for Extreme Day Classification in Electricity Markets

Saud Alghumayjan, Ming Yi, Bolun Xu
Dept. of Earth and Environmental Engineering
Columbia University
New York, NY 10027, USA
{saa2244, my2826, bx2177}@columbia.edu

Abstract—This paper proposes a few-shot classification framework based on Large Language Models (LLMs) to predict whether the next day will have spikes in real-time electricity prices. The approach aggregates system state information, including electricity demand, renewable generation, weather forecasts, and recent electricity prices, into a set of statistical features that are formatted as natural-language prompts and fed to an LLM along with general instructions. The model then determines the likelihood that the next day would be a spike day and reports a confidence score. Using historical data from the Texas electricity market, we demonstrate that this few-shot approach achieves performance comparable to supervised machine learning models, such as Support Vector Machines and XGBoost, and outperforms the latter two when limited historical data are available. These findings highlight the potential of LLMs as a data-efficient tool for classifying electricity price spikes in settings with scarce data.

Index Terms—Electricity markets, Machine learning, Large Language Models

I. INTRODUCTION

The short-term behavior of electricity prices has become increasingly difficult to predict as modern power systems incorporate higher levels of renewable generation and face greater variability in operating conditions. These have led to more frequent fluctuations in net load and produced highly volatile real-time price patterns [1]. Because operators and market participants rely on accurate next-day forecasts for scheduling, hedging, and risk management, the ability to anticipate price spikes is increasingly important. In practice, many market participants rely on automated bidding and risk-management software trained on the majority of normal days. While these systems perform reliably under typical conditions, they often perform poorly during extreme price events, leading to financial losses and inefficient market responses. However, forecasting these events remains challenging due to the combined effects of operational uncertainty, renewable forecast errors, evolving demand profiles, adverse weather conditions, and persistent transmission congestion.

A wide range of statistical and machine-learning models has been applied to electricity price forecasting, with most focusing on regression-based prediction. Traditional time-series methods and modern deep-learning architectures, including

Multi-Layer Perceptron (MLP) [2], hybrid neural networks [3], and transformer-based models [4], aim to directly estimate future prices. Conformal prediction has also been explored to quantify predictive uncertainty [5]. While these methods capture general price dynamics, they struggle with rare and highly nonlinear spike events, which are overshadowed by the predominance of non-spike prices in training.

In practical operations, market participants frequently need to know whether the next day is likely to experience abnormally high prices, the ability to flag such days in advance would warrant hedging or operational adjustments. This motivates electricity spike classification, in which the goal is to determine whether prices will exceed a critical threshold. Prior studies have investigated such approaches using support vector machines (SVMs) [6] and weighted K-nearest neighbor (WKNN) models [7], with additional methods surveyed in [8]. Although these approaches improve sensitivity to extreme events, they typically require large labeled datasets and may degrade when spikes are infrequent or when system conditions shift over time. Such spike days also affect residential consumers, who are generally risk-averse. With the increasing adoption of virtual power plants that aggregate residential batteries for market participation [9], poor anticipation of these days can lead to inefficient bidding and financial losses.

Large language models (LLMs) can interpret structured numerical data expressed in natural language and link system signals, including demand variability, renewable output, and price trends, to price outcomes. They also perform well with limited labeled data, addressing the rarity of spike events and changing power conditions. Recently, LLMs has been used in some power system applications [10]–[13]. Motivated by this, we propose a few-shot LLM-based classifier that estimates the likelihood of next-day real-time price spikes from textual system descriptions. Our main contributions are:

- We develop a few-shot classification framework that leverages LLMs to flag whether the next day will be an electricity price spike day in the real-time market, based on statistical summaries of system conditions.
- We introduce a simple but effective procedure that converts market and system variables, including forecasts and recent price dynamics, into textual features suitable for LLM-based inference.
- We evaluate the proposed framework using historical

The authors would like to acknowledge supports from Columbia Data Science Institute, Red River Clean Energy, King Abdulaziz City for Science and Technology (KACST), and National Science Foundation under award ECCS-2239046.

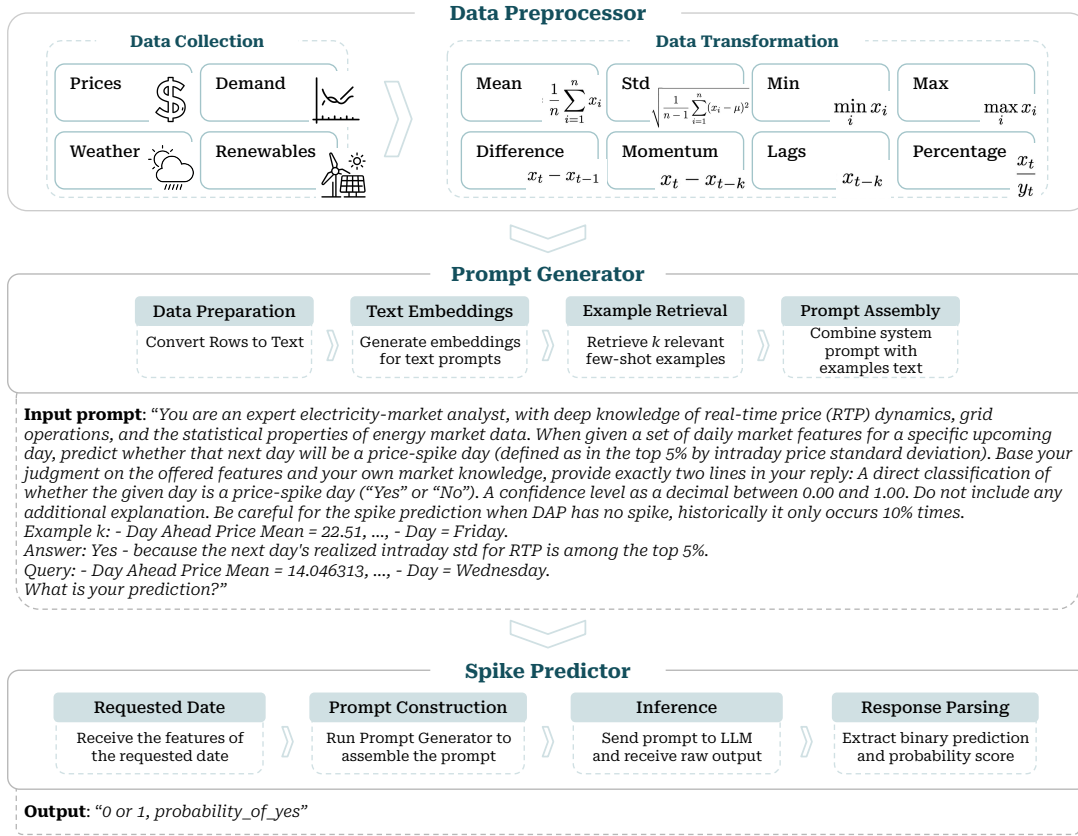


Fig. 1. The pipeline for the proposed approach.

Texas market data and benchmark its performance against SVM and XGBoost, despite the latter two being trained on a much larger dataset.

- We demonstrate that LLM-based predictors are effective in limited-data settings, highlighting their value as a data-efficient approach to electricity price classification.

The remainder of the paper is organized as follows: Section II introduces the framework. Section III presents the case study, and Section IV concludes the paper.

II. LLM-BASED SPIKE CLASSIFIER FRAMEWORK

We present an LLM-based framework to predict whether the next day will feature real-time electricity price spikes. The framework consists of three sequential components: a *data preprocessor*, which integrates and transforms electricity system wide datasets into structured learning features; a *prompt generator*, which converts daily feature vectors into textual representations and retrieves relevant few-shot examples using embedding-based similarity search; and a *spike predictor*, which assembles the final prompt, queries the LLM, and parses the resulting prediction. These components form a complete pipeline, spanning from raw system information to spike-day probability forecasts, without relying on model training. Figure 1 shows a high-level diagram of the full workflow.

A. Data Preprocessor

The first component of the framework collects all the information that describes the state of the electricity system for each day. These variables are collected from several sources, including day-ahead and real-time market prices, forecasts and actual net demand, forecasted reserve levels, renewable generation forecasts, and weather data. Each dataset is aligned and transformed into a collection of engineered features intended to summarize system conditions. These include daily statistics such as means and standard deviations, lagged values that incorporate information from the previous day, and multi-day summaries that capture recent trends. Weather-related measures, renewable penetration, and calendar indicators are also included to account for environmental and temporal patterns. We define a spike day as any day whose real-time intraday price standard deviation exceeds the 95% threshold, where

$$\text{Spike}_t = \begin{cases} 1, & \text{if } \sigma_{\text{intraday},t}^{\text{RTP}} > \text{Quantile}_{0.95}(\sigma_{\text{intraday}}^{\text{RTP}}), \\ 0, & \text{otherwise,} \end{cases}$$

Table I provides an overview of all engineered features, categorized by their underlying data source. This feature set forms the daily feature vector that is later translated into text for use in the prompt generation stage.

TABLE I
ENGINEERED FEATURES BY DATA CATEGORY

Category	Feature
Day-Ahead Price	dap_mean
	dap_std
	dap_max
	dap_min
	dap_momentum
	dap_mean_3d dap_std_3d
Real-Time Price	rtp_mean_lag1
	rtp_std_lag1
Net Demand (Forecast)	net_demand_forecast_mean
	net_demand_forecast_min
	net_demand_forecast_max
	net_demand_forecast_std
	net_demand_diff
Net Demand (Actual)	net_demand_actual_mean_lag1
	net_demand_actual_std_lag1
Forecasted Reserves	reserve_mean
	reserve_std
	reserve_max
	reserve_min
Weather	temp_max
	temp_min
	temp_diff
	wind_speed_max
	wind_speed_min
Renewables	pct_renewable_forecast
Calendar	day
	month
Labels	is_dap_spike
	is_spike

B. Prompt Generator

The second component of the framework converts each day's feature vector into natural-language descriptions suitable for use in large language models. For each row, the feature values are transformed into a textual prompt. These text prompts are then mapped into numerical embedding using an embedding model. To efficiently locate past days that resemble the query day, we use FAISS [14], a library designed for fast similarity search in large collections of vectors.

FAISS builds an index over the embeddings that enables quick nearest-neighbor queries, even when the dataset is large. Using this index, the prompt generator identifies a set of candidate examples that are most similar to the day being predicted. When predicting for a query date, the prompt generator retrieves the most similar historical spike day and then selects the closest days to form a candidate pool. From these candidates, a set of k examples is selected using maximal marginal relevance (MMR) [15], which helps choose examples that are relevant to the query and not redundant with one another. The final prompt provided to the LLM consists of the system instruction, the selected examples with their labels, and the text description of the query day. The full procedure is summarized in Algorithm 1.

Algorithm 1: Prompt Generator

Data: Input dataframe $\mathbf{D}$
1 Require: System prompt S , number of examples k , MMR parameter λ , embedding model $\mathcal{E}$
2 Data Preparation
3 foreach row r **in** $\mathbf{D}$ **do**
 4 Convert feature values of r into a textual description t_r ;
5 end
6 Text Embeddings
7 foreach text prompt t_r **do**
 8 Compute embedding e_r using model $\mathcal{E}$;
9 end
10 Store embeddings $\{e_r\}$;
11 Example Retrieval
12 Normalize embeddings and construct FAISS index using $\{e_r\}$;
13 Convert query row to text prompt t_q and compute embedding e_q ;
14 Identify valid prior rows $r < q$;
15 Select the most similar prior spike example;
16 Retrieve additional nearest-neighbor candidates from the FAISS index;
17 Apply Maximal Marginal Relevance (MMR):
18 Initialize selected set $\mathcal{S}$;
19 while $|\mathcal{S}| < k$ **do**
 20 Compute MMR score for each candidate using query similarity and redundancy penalty λ ;
 21 Add the highest-scoring candidate to $\mathcal{S}$;
22 end
23 Prompt Assembly
24 foreach example $e_i \in \mathcal{S}$ **do**
 25 Format example text and associated label;
26 end
27 Append the query description;
28 Concatenate system prompt S , retrieved examples, and query description;
29 return Final prompt ready for LLM inference.

C. Spike Predictor

The final component executes LLM-based inference. Given a requested date, the spike predictor obtains the corresponding feature vector, calls the prompt generator to construct the full textual prompt, and sends this prompt to the LLM. The LLM returns a structured two-line response: a binary classification (Yes/No) indicating whether the next day is expected to be a spike day, and a confidence score between 0 and 1. The predictor parses this output to produce the final prediction and associated probability.

III. CASE STUDY

This section evaluates the proposed LLM-based spike classification framework using historical data from the Texas electricity market. We describe the experimental setup, the




Fig. 2. Prediction probabilities for XGBoost, SVM, and LLM models, with PR and ROC thresholds shown for comparison.

classification metrics used for performance assessment, and the results obtained across multiple thresholding strategies and data availability scenarios.

A. Experiment Setup

We utilize the Electric Reliability Council of Texas (ERCOT) system-wide data covering the period from 2021 to 2024. The years 2021-2023 are used for training and 2024 is used for testing. All models produce a one-day-ahead prediction of whether the next day will be a spike day. The engineered features described in Section II serve as inputs to the prompt generator. For the LLM-based classifier, we use the `gpt-4.1` model [16]. Benchmark models include SVM [17] and XGBoost [18] classifiers, both trained on the same feature set. All computations were carried out on a high-performance computing cluster with Intel Xeon Platinum 8640Y 2Ghz CPUs, 512GB memory, and 2xNVIDIA A40 GPUs.

B. Classification Metrics

To evaluate the performance of the spike classifiers, we rely on standard measures derived from the confusion matrix. Let TP, FP, FN, and TN denote the numbers of true positives, false positives, false negatives, and true negatives:

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP}) \quad (1)$$

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN}) \quad (2)$$

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{FP} + \text{FN} + \text{TN}) \quad (3)$$

$$\text{F1-score} = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (4)$$

Precision reflects the fraction of predicted spikes that are correct, while recall measures the fraction of actual spike days correctly identified. Accuracy provides the overall proportion of correct predictions, and the F1-score summarizes the balance between precision and recall.

In addition to these point metrics, we examine the receiver operating characteristic (ROC) curve and the precision–recall (PR) curve. The ROC curve plots the true positive rate against the false positive rate as the classification threshold varies, highlighting the classifier’s ability to separate spike and non-spike days across all thresholds.

The PR curve plots precision against recall across thresholds. This curve informs settings where rare positive events occur. Because spike days constitute a small fraction of the dataset, precision–recall analysis provides a better understanding of the classifier behavior. We examine both ROC and PR curves and derive threshold values from each: the ROC-optimal threshold balances sensitivity and false-alarm rate, and the PR-optimal threshold maximizes the F1-score.

C. Results

Table II summarizes the classification outcomes under the three thresholding settings. When using the full training set, the LLM performs competitively across several metrics despite not being trained directly on historical labels. In particular, the LLM often achieves the highest or second-highest among all metrics under multiple thresholding scenarios. These results indicate that the LLM can effectively leverage prompt-generated context to detect spike events, comparable to supervised learning models that rely on extensive training data.

Figure 2 plots the predicted spike probabilities for each model across the test period, along with the PR-optimal and ROC-optimal thresholds. The LLM produces smoother probabilities and tends to concentrate high-confidence predictions around periods of spikes, while the SVM and XGBoost models exhibit more fluctuation that reflect less confident on predictions. This explains why the LLM model has higher thresholds than the other two models.

TABLE II
MODEL PERFORMANCE ACROSS THRESHOLD TYPES

Model	TP	FP	FN	TN	Prec.	Rec.	Acc.	F1
Fixed (0.5)								
Thresholds: XGBoost = 0.50; SVM = 0.50; LLM = 0.50;								
XGBoost	0.0	3.0	9.0	354.0	0.0	0.0	96.72	0.0
SVM	1.0	0.0	8.0	357.0	100.0	<u>11.11</u>	97.81	<u>20.0</u>
LLM	2.0	3.0	7.0	354.0	<u>40.0</u>	22.22	<u>97.27</u>	28.57
ROC-optimal								
Thresholds: XGBoost = 0.04; SVM = 0.04; LLM = 0.13;								
XGBoost	6.0	43.0	3.0	314.0	<u>12.24</u>	<u>66.67</u>	<u>87.43</u>	<u>20.69</u>
SVM	9.0	140.0	0.0	217.0	6.04	100.0	61.75	11.39
LLM	6.0	37.0	3.0	320.0	13.95	<u>66.67</u>	89.07	23.08
PR-optimal								
Thresholds: XGBoost = 0.17; SVM = 0.29; LLM = 0.91;								
XGBoost	4.0	11.0	5.0	346.0	26.67	44.44	95.63	<u>33.33</u>
SVM	3.0	5.0	6.0	352.0	<u>37.5</u>	<u>33.33</u>	<u>96.99</u>	35.29
LLM	2.0	2.0	7.0	355.0	50.0	22.22	97.54	30.77

TABLE III
MODEL PERFORMANCE ACROSS THRESHOLD TYPES - LIMITED DATA

Model	TP	FP	FN	TN	Prec.	Rec.	Acc.	F1
Fixed (0.5)								
Thresholds: XGBoost = 0.50; SVM = 0.50; LLM = 0.50;								
XGBoost	1.0	74.0	8.0	283.0	1.33	<u>11.11</u>	77.6	2.38
SVM	0.0	0.0	9.0	357.0	0.0	0.0	97.54	0.0
LLM	3.0	7.0	6.0	350.0	30.0	33.33	<u>96.45</u>	31.58
ROC-optimal								
Thresholds: XGBoost = 1.00; SVM = 0.03; LLM = 0.93;								
XGBoost	0.0	0.0	9.0	357.0	0.0	0.0	97.54	0.0
SVM	2.0	75.0	7.0	282.0	2.6	<u>22.22</u>	77.6	4.65
LLM	3.0	5.0	6.0	352.0	37.5	33.33	<u>96.99</u>	35.29
PR-optimal								
Thresholds: XGBoost = 0.03; SVM = 0.03; LLM = 0.93;								
XGBoost	9.0	357.0	0.0	0.0	2.46	100.0	2.46	4.8
SVM	2.0	75.0	7.0	282.0	2.6	22.22	<u>77.6</u>	4.65
LLM	3.0	5.0	6.0	352.0	37.5	<u>33.33</u>	96.99	35.29

Finally, Table III reports performance in a limited-data setting, where only two months of historical data are used for training, instead of the full three years. In this scenario, the LLM considerably outperforms the supervised baselines across most metrics. Whereas SVM and XGBoost suffer substantial degradation when training data are scarce, the LLM remains stable. This demonstrates a key advantage of the proposed approach when labeled training data are highly limited.

IV. CONCLUSION

This paper introduces an LLM-based framework for predicting whether the next day will be a day of spikes in real-time electricity prices using a few-shot classification approach. The method integrates system-level data, engineered features, and a structured prompt design to enable large language models to perform spike classification without direct model training. Through a combination of textual feature representation, similarity-based example retrieval, and LLM inference, the proposed approach achieves performance comparable to established machine learning models despite having no su-

pervised training component. Our results demonstrate that, while traditional models such as XGBoost and SVM perform well when large training datasets are available, their performance deteriorates substantially under limited-data conditions, during which the LLM-based classifier remains competitive, highlighting its value as a data-efficient classification tool. For the next steps, several extensions are possible, such as expanding the prompt format to include multi-day contextual information or fine-tuning domain-specific LLMs, which may further enhance spike prediction accuracy.

REFERENCES

- [1] Y. Wang, Y. Dvorkin, R. Fernández-Blanco, B. Xu, and D. S. Kirschen, "Impact of local transmission congestion on energy storage arbitrage opportunities," in *2017 IEEE Power & Energy Society General Meeting*. IEEE, 2017, pp. 1–5.
- [2] I. P. Panapakidis and A. S. Dagoumas, "Day-ahead electricity price forecasting via the application of artificial neural network based models," *Applied Energy*, vol. 172, pp. 132–151, 2016.
- [3] M. Lehna, F. Scheller, and H. Herwartz, "Forecasting day-ahead electricity prices: A comparison of time series and neural network models taking external regressors into account," *Energy Economics*, vol. 106, p. 105742, 2022.
- [4] S. Alghumayjan, J. Han, N. Zheng, M. Yi, and B. Xu, "Energy storage arbitrage in two-settlement markets: A transformer-based approach," in *Proceedings of the Power System Computation Conference (PSCC) 2024*, Paris, France, 2024.
- [5] S. Alghumayjan, B. Xu, and M. Yi, "Conformal uncertainty quantification of electricity price predictions for risk-averse storage arbitrage," in *2025 IEEE Power Energy Society General Meeting (PESGM)*, 2025, pp. 1–5.
- [6] C. Zhang, Y. Fu, and L. Gong, "Short-term electricity price forecast using frequency analysis and price spikes oversampling," *IEEE Transactions on Power Systems*, vol. 38, no. 5, pp. 4739–4751, 2022.
- [7] S. Liu, Y. Jiang, Z. Lin, F. Wen, Y. Ding, and L. Yang, "Data-driven two-step day-ahead electricity price forecasting considering price spikes," *Journal of Modern Power Systems and Clean Energy*, vol. 11, no. 2, pp. 523–533, 2022.
- [8] J. Nowotarski and R. Weron, "Recent advances in electricity price forecasting: A review of probabilistic forecasting," *Renewable and Sustainable Energy Reviews*, vol. 81, pp. 1548–1568, 2018.
- [9] S. Hwang and N. Kittner, "Virtual power plants: bridging utility and customer gaps to facilitate clean energy transitions," *Progress in Energy*, vol. 7, no. 2, p. 022008, 2025.
- [10] X. Lu, J. Qiu, Y. Yang, C. Zhang, J. Lin, and S. An, "Large language model-based bidding behavior agent and market sentiment agent-assisted electricity price prediction," *IEEE Transactions on Energy Markets, Policy and Regulation*, 2024.
- [11] Y. Zhou and M. Wang, "Empower pre-trained large language models for building-level load forecasting," *IEEE Transactions on Power Systems*, 2025.
- [12] A. Jena, F. Ding, J. Wang, Y. Yao, and L. Xie, "Llm-based adaptive distribution voltage regulation under frequent topology changes: An in-context mpc framework," *IEEE Transactions on Smart Grid*, 2025.
- [13] M. Jia, Z. Cui, and G. Hug, "Enhancing llms for power system simulations: A feedback-driven multi-agent framework," *IEEE Transactions on Smart Grid*, 2025.
- [14] M. Douze, A. Guzhva, C. Deng, J. Johnson, G. Szilvasy, P.-E. Mazaré, M. Lomeli, L. Hosseini, and H. Jégou, "The faiss library," 2024.
- [15] J. Carbonell and J. Goldstein, "The use of mmr, diversity-based reranking for reordering documents and producing summaries," in *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, 1998, pp. 335–336.
- [16] OpenAI, "Gpt-4.1," 2024, large language model. [Online]. Available: <https://platform.openai.com/>
- [17] N. Cristianini and E. Ricci, "Support vector machines," in *Encyclopedia of algorithms*. Springer, 2008, pp. 928–932.
- [18] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.