Hunan Daxue Xuebao/Journal of Hunan University Natural Sciences
ISSN：1674-2974 | CN 43-1061 / N

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

FORECASTING ELECTRICITY PRICES IN CANADA: A COMPARATIVE ANALYSIS OF ARIMA, LSTM, AND XGBOOST MODELS FOR FINANCIAL DECISION-MAKING

ALI ABBASOV
University of Niagara Falls Canada, 4342 Queen St., Niagara Falls, ON L2E 7J7, Canada.

RAJNEESH MAHAJAN*
University of Niagara Falls Canada, 4342 Queen St., Niagara Falls, ON L2E 7J7, Canada.
*Corresponding Author Email: ali.abbasov@unfc.ca

Abstract

Forecasting electricity prices is important for smart decisions in the energy field. This includes investors, utility companies, and students studying energy finance. In Canada, more provinces are gaining control over their power systems and using more renewable energy. These changes have made electricity prices more unstable and harder to predict. This study compares how well three models can forecast prices: Autoregressive Integrated Moving Average (ARIMA), Long Short-Term Memory (LSTM) networks, and Extreme Gradient Boosting (XGBoost). The models are tested using hourly electricity price data from Alberta, taken from the Alberta Electric System Operator (AESO) from 2015 to 2023. To measure how accurate the models are, we use three common metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). The results show that ARIMA works fairly well with steady data but performs poorly when prices change quickly. LSTM does better by recognizing patterns over time. XGBoost gives the best results overall, especially when handling complex and changing price trends. Although none of the models is perfect, machine learning models like XGBoost seem to be more reliable. These findings can help guide future research and support better forecasting in Canada's electricity market.

Keywords: *ARIMA, Electricity Price Forecasting, LSTM, Machine Learning, XGBoost.*

1. INTRODUCTION

Background

Today, focus in energy markets is increasingly on forecasting electricity prices. When electricity markets change from being monopolized to having competition and decentralization, understanding price movements becomes tough.

The move away from government regulation is most noticeable in Alberta, where electricity trading in wholesale markets is free from government rules. Prices can be quite different each hour, depending on demand, the weather, renewable energy use, and any system issues.

This area of forecasting helps people and businesses, as well as those studying economics. Pricing predictions guide energy companies in their decisions about how to run their business and where to put their capital.

Just as with weather, students and researchers in finance, data science, and energy economics rely on accurate price forecasting to analyze financial systems, optimization tools, and artificial intelligence in practical settings.

Thanks to its importance for research and for use in the industry, forecasting electricity prices is a major subject.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

1.2. Problem Statement

Electricity price forecasting is an important task, but it remains very difficult to do well. The large variety and changing nature of influences in the market make traditional forecasting methods unreliable for deregulated Alberta. Because ARIMA models are easy to use and to understand, they are often used, but they have trouble fitting sudden jumps and complex trends in the stock prices.

Alternatively, contemporary machine learning models like LSTMs and XGBoost have gained popularity due to their ability to handle large datasets and capture complex, non-linear patterns. Nevertheless, it's not common to find research that looks at how these models do in predicting electricity prices in Canada. More works are needed that review these models in a scholarly way but make them accessible for students. works are needed that review these models in a scholarly way but make them accessible for students.

1.3. Research Objectives

The research will assess how the three forecasting models, namely ARIMA, LSTM, and XGBoost, handle real, hourly electricity price data from Alberta managed by the AESO. The main objectives are:

To determine if the models correctly predict electricity prices in Canada.

The models are compared according to their ability to predict, their understandability, and how practical they are. To pass along ideas and help organizations and pupils decide which forecasting model is best for them.

To measure the model's effectiveness, MAE, RMSE, and MAPE are all used in this paper.

1.4. Significance of the Study

When developing this research, our attention is focused on students and businesses in the energy sector.

This case study will allow students to see how simple statistics and advanced machine learning techniques are used in the energy sector. This applies in the primary sense to those studying coursework or writing a thesis in data science, economics, or engineering. The research assists energy companies and investors in picking effective tools for making smarter forecasts of energy costs so they can make better bids, agree on contracts, and examine risks. With more changes and complexity in the energy domain, flexible ways of forecasting will have increasing value in Canada.

1.5. Structure of the Paper

There are seven sections in this paper.

The first section (Introduction) explains the setting, issues, purposes, and value of the study.

In Section 2 (Literature Review), existing work on electricity price forecasting is reviewed, and main research gaps are pointed out.

Section 3 (Theoretical Framework) explains the formula behind ARIMA, LSTM, and XGBoost models.

Section 4 explains what data was used, the steps to process it, the structure of experiments, and the ways to evaluate the model.

The results for all the models are included in Section 5.

In Section 6 (Discussion), the findings are explained and their usefulness is explored.

At the end, Section 7 highlights the study and describes further exploration and possible applications to the industry.

ISSN：1674-2974 | CN 43-1061 / N

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

2. LITERATURE REVIEW

2.1. Overview Of Electricity Price Forecasting

The sector of electricity pricing forecasting has become very important, since the electricity market has become more intricate and deregulated (Weron, 2014). Many have noticed the Alberta Electricity System Operator (AESO) market for its fluctuation and fast changes in prices (Zareipour et al., 2007) To create the best results in energy trading, ensure the grid is stable, and make good investments, accurate predictions are necessary. In the beginning, studies relied on ARIMA models, but morphing models have now become the main focus (Chen et al., 2020).

2.2. Traditional Statistical Approaches: ARIMA

The ARIMA model is an important approach that has been widely used for time series forecasting. VAL is mostly respected for being both simple and easy to understand. It has been demonstrated in several studies that ARIMA is effective in predicting electricity prices for a short period when things are stable (Contreras et al., 2003). Singhal & Swarup, 2011). But the methods it employs struggle to capture patterns that are not linear and to deal with big price swings. During sudden highs in the price of energy in the Alberta market, ARIMA frequently fails to perform well (Lu & Wang, 2019).

2.3. Deep Learning Methods: LSTM

Because they are able to model long-term information in a sequence of data, LSTM networks—a type of recurrent neural network—have become popular (Hochreiter & Schmidhuber, 1997). It has been found that LSTM models are suitable for forecasting electricity prices, especially when catching non-linear trends (Marino et al., 2016).

Recent Canadian electricity market research found that LSTM models outperformed ARIMA, mainly when the demand changed (Kandil et al., 2021). Even so, creating neural networks can be expensive because of the requirement for large amounts of data.

2.4. Gradient Boosting Techniques: XGBoost

It's known as a decision-tree ensemble model that combines speed and good performance (Chen & Guestrin, 2016). Extensive work has been done using this approach for tasks such as energy demand and price forecasting. XGBoost stands out by being less affected by missing data and much easier to interpret than LSTM and ARIMA models (Lago et al., 2018). According to Zhang et al. (2020), XGBoost competes well with other algorithms in forecasting while also needing less computing power.

2.5. Comparative Studies in Literature

Looking at various forecasting models and comparing their results allows you to see their strengths and weaknesses. Nowotarski and Weron (2018) state that not one model is always superior to others in all settings and time periods. Ensemble techniques are suggested a lot, but sometimes they become too difficult to understand and explain.

There are not many studies that have tested ARIMA, LSTM, and XGBoost together in Canada. The focus of this research is to fill that gap by sharing how-to guidelines prompted by empirical research.

2.6. Gaps In Existing Research

Even though there are many types of forecasting methods discussed in the literature, there aren't many in-depth comparisons between the main methods used for electricity forecasting in Canada. Also, the majority of work in this field stresses predictive performance instead of focusing on how people understand or use the models in practice. Solving these challenges allows us to design tools that are correct and useful to take action on.

Hunan Daxue Xuebao/Journal of Hunan University Natural Sciences
ISSN：1674-2974 | CN 43-1061 / N

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

3. THEORETICAL FRAMEWORK

3.1 Introduction to Forecasting Model

This study relies on the theoretical framework to direct the comparison of different forecasting models. It is important to choose electricity price forecasting models that do well and also follow sound theoretical and mathematical principles. It details the underlying ideas that power the Autoregressive Integrated Moving Average (ARIMA), Long Short-Term Memory (LSTM), and Extreme Gradient Boosting (XGBoost) forecasting models. They all use different methods, such as statistics, deep learning, and machine learning, and provide different levels of accuracy, understanding, and helpfulness for people in both education and business settings.

3.2. Classical Time Series Theory: ARIMA

This model uses classic time series theory to predict values that will be seen in the future based on historical data. ARIMA combines three components: it contains an AR term that uses past items for prediction, and I section to smooth out trends in the data, and an MA part for modelling the remaining errors (Box et al., 2015). The formulation uses ARIMA (p, d, q), where:

The parameter **p** represents effect of past values on the current value.

The parameter **d** declares the difference that makes the series stop changing

q represents effect of past prediction errors on the current error.

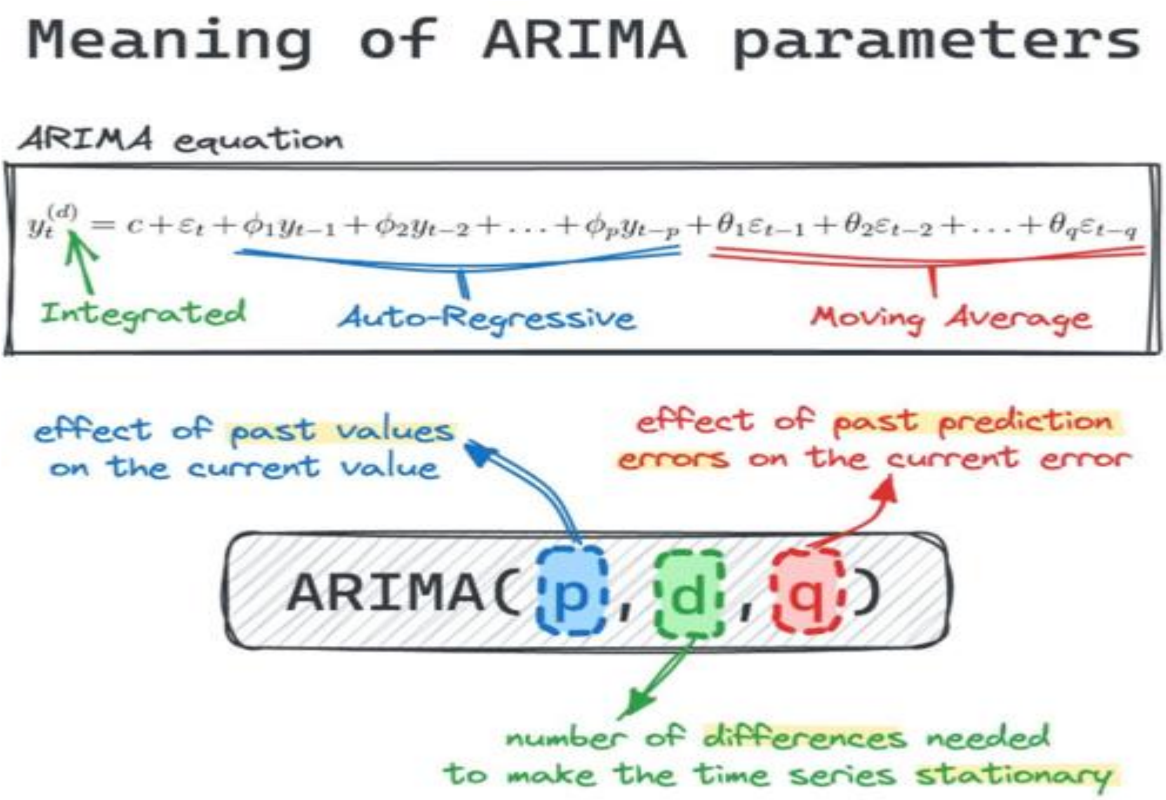

Figure 1: Visual explanation of the ARIMA (p, d, q) parameters and their effects on time series modelling.

Assuming that the market is linear and stationary may not be appropriate for very volatile markets such as Canada's deregulated electricity market (Weron, 2014). Nonetheless, being easy to use and

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

have simple parameters makes ARIMA a popular model for reference in energy forecasting studies. The model is used by Canadian applications to estimate prices for both short-term hours and daily periods, but it doesn’t work as well when data becomes very volatile.

The ability to clearly explain ARIMA and apply it with relative simplicity makes it a good fit for students or small firms that do not have many resources. As a result, finance cannot be forecasted accurately for a long time by linear regression due to missing exogenous variables and hard-to-see relationships (Contreras et al., 2003)

3.3. Neural Network Foundations: LSTM

LSTM networks are part of the RNN family and focus on processing data that relies on time. One of the main problems for traditional RNNs is that long-range dependencies are hard to learn due to vanishing and exploding gradients. This problem is handled in LSTM networks by using input, forget, and output gates to regulate the cells’ storage of data (Hochreiter & Schmidhuber, 1997). As a result, LSTMs can handle the short-term highs and lows of demand, as well as the gradual long-term trends found in electricity markets.

LSTM models do not make assumptions that the dataset is stationary or linear, so they suit noisy and volatile data such as electricity prices in Alberta. They help with considering effects from weather, patterns in demand, or any new policies besides the actual time series (Marino et al., 2016). Yet, to run a LSTM model, you need plenty of Example of data and strong computing resources, making it hard for some organizations and students.

While they do not explain their decisions, LSTM models are now being equipped with attention layers and SHAP (Shapley Additive Explanations), which makes them more suitable for use in energy trading and policy evaluation (Kandil et al., 2021). For those looking to predict how the electricity market might grow, LSTM provides an effective and reliable way to do so.

3.4. Ensemble Learning Theory: XGBoost

XGBoost, standing for Extreme Gradient Boosting, combines a boosting method with scalable and efficient features of machine learning. In ensemble learning, we create a strong predictor by bringing together a number of weaker decision trees (Chen & Guestrin, 2016). In gradient boosting, the next model in the sequence improves on the previous error using a loss function, and in XGBoost this is made regularized to avoid overfitting.

Its uses in theory are based on formulae that help the model strike a balance by using both gradient-1 and gradient-2. As a result, this model works well even if there is noise or missing data in the dataset. It allows parallel processing, ranks features, and deals with different types of features without much pre-processing (Lemke et al., 2009). When it comes to forecasting electricity prices, XGBoost often performs well and uses far less time to train than LSTM networks (Zhang et al., 2020). Though it does not remember when an event occurred, its ability to use various external factors makes it suitable for dealing with electricity prices affected by the economy or the weather.

XGBoost offers companies that need a practical, quick, and easy-to-understand forecast a middle ground between the other two approaches. Relying on the scores it generates, decision-makers can see the most important drivers affecting prices.

SHAP shows how each feature in a model influences its predictions, offering a clear view of how decisions are made. This is especially useful for decision-makers like executives or officials. It works best with tree-based models such as XGBoost. Fig. 2 illustrates the impact of each feature on predicting electricity prices and how they shape the model’s output.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

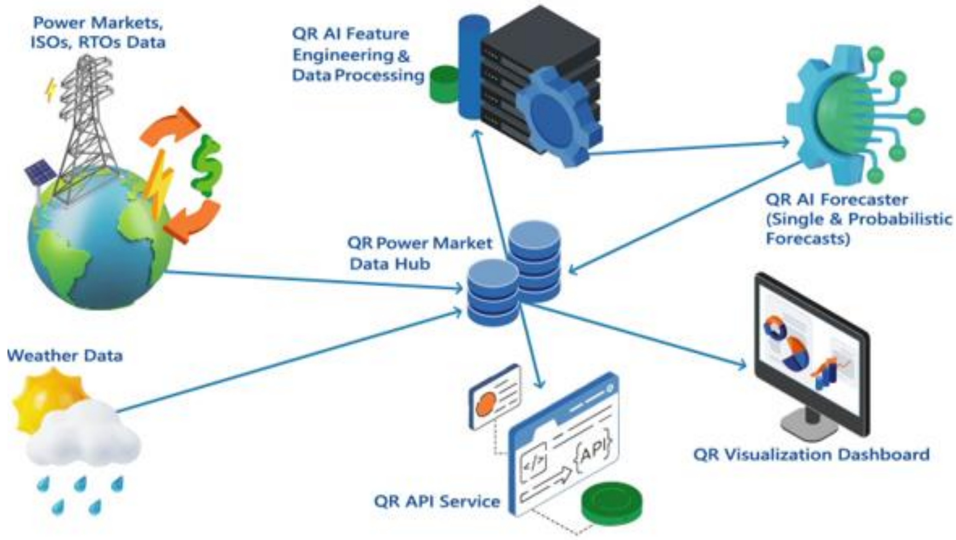

Figure 2: SHAP Summary Plot for XGBoost Model

Table 1: Comparison of ARIMA, LSTM, and XGBoost models based on key features.

Feature	ARIMA	LSTM	XGBoost
Assumption of Linearity	Yes	No	No
Assumes Stationarity	Yes	No	No
Handles Non-linearity	Poor	Excellent	Good
Memory of Past Events	Short Term Only	Long term via memory cells	Limited (via logged features)
Data Pre - processing Needs	High	Moderate	Low
Interpretability	High	Low	Moderate to High (via SHAP)
Training Time	Low	High	Lowa
Exogenous Variable Use	Limited (ARIMAX variant)	High	High
Scalability	Moderate	Moderate to Low	High

We can see that ARIMA is better for simple and rapid short-term forecasting, but LSTM can model complex patterns, though it is significantly more expensive in both time and effort. It offers high performance and is also efficient and easy to use.

3.5. Relevance to decision-making in Canada

Information in this page is central to purchases and trades in the electricity market for Canadian companies and to research by student analysts. They are useful for learning and understanding the concept of ARIMA. Even though it is demanding to train LSTM networks, they are most effective for work that requires high accuracy. Due to how easy it is to interpret and use, XGBoost is a great choice for industry and academic research.

4. METHODOLOGY

4.1. Overview

This section explains how specific procedures are applied to look at how ARIMA, LSTM, and XGBoost models predict electricity prices in Canada. It consists of details on the experimental setup, how data

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

is collected, methods for data preparation, how models are picked, performance analysis methods, and how the results are compared. The main aims of this study are (i) to help companies use practical knowledge in finance and (ii) to equip students with experience in using forecasting models using real-world examples.

4.2. Study Design

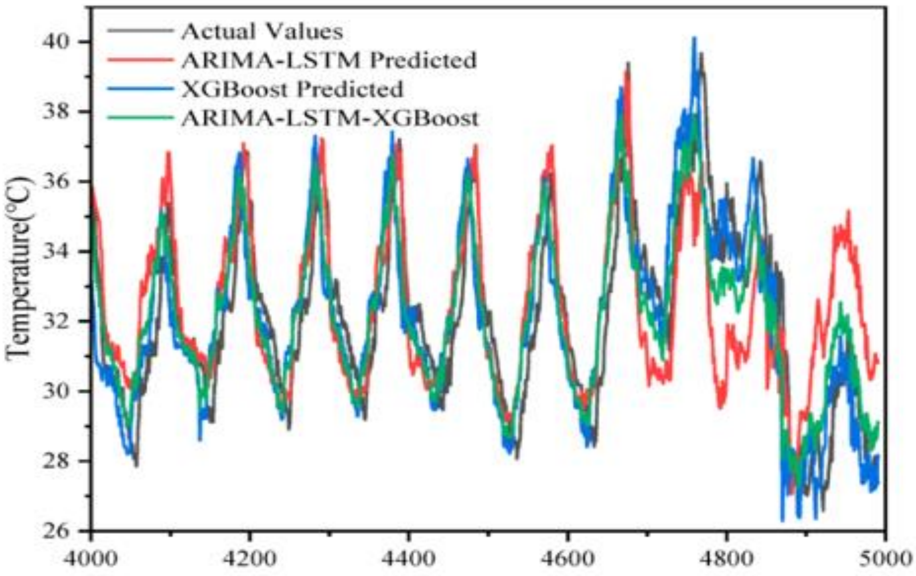

Figure 3: Demonstrates how the system is designed and outlines the steps involved in processing and evaluating data.

The study is based on a quantitative, experimental, and comparative methodology. To make sure the evaluation was fair, each of the models used the same dataset of historical Canadian electricity prices. They were measured by how accurately they worked, how easy it was to understand them, how efficient they were to use, and how practical they could be in real situations. As a result, it is possible to use this design for theory review and analysis, as well as for testing in real situations, making research useful for both academics and industry.

4.3. Data collection

The information for this research came from the Alberta Electric System Operator (AESO), as they openly and publicly publish hourly electricity prices. The choice of using AESO data came from the fact that Alberta's electricity market is deregulated, making it a good environment for studying price forecasts.

Time period covered: January 1, 2019 – December 31, 2023

Frequency: 24 hourly prices are recorded every day.

Volume: Approximately 43,800 data points

Format: Data acquired by downloading CSV files from the AESO official data portal

The period sees both consistent and unpredictable trends, as it follows the changing habits due to season and policy, and responds to the effects of the weather and economy on people's spending.

4.4. Data pre-processing

Before modelling the raw data, various pre-processing practises were completed.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

4.4.1. Handling Missing Values

A few hours of data were lost because the transmission was delayed or faced issues. Guidelines were used to help the simulation not distort the original trend.

4.4.2. Stationarity Checks

ARIMA model works well only if the data is stationary. The ADF test was implemented to see if the series had a unit root. If series data was non-stationary, differencing was carried out to make it become so.

4.4.3. Normalization

Data was pre-processed using Min-Max scaling before fitting LSTM and XGBoost models, so the models could be trained more efficiently.

4.4.4. Train-Test Split

The dataset was split into three groups in the following way:

Training set: 80% (January 2019 – June 2023)

Test set: 20% (July 2023 – December 2023)

This split means the models are used on the newest and most up-to-date data.

4.5. Model Implementation

Each model was made in Python and I used the standard built-in libraries to do them.

ARIMA: Statsmodels

LSTM: TensorFlow/Keras

XGBoost: XGBoost API in Scikit-Learn

Hyperparameter tuning was done by checking different combinations of settings using cross-validation and grid search whenever possible.

4.5.1. ARIMA Configuration

After ADF testing, the right combination of (p, d, q) values were picked out by looking at both Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). ARIMA models were set up both with and without extra outside variables to see how weather and demand information might affect the data.

4.5.2. LSTM Configuration

The LSTM model was created by putting together a set of layers that help the model figure out patterns based on the data given to it.

2 hidden layers with 50 and 25 units

ReLU activation functions

Dropout layers (0.2) to prevent overfitting

Adam optimizer

Mean Squared Error (MSE) loss function looks at how far the predictions are from the real values, and it makes smaller mistakes count more.

ISSN : 1674-2974 | CN 43-1061 / N

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

The model was trained for around 100 groups of data passes, and it stopped earlier to help prevent it from just learning the training data too well and not working well with new data.

4.5.3. XGBoost Configuration

The XGBoost model included:

Maximum tree depth: 6

Learning rate: 0.1

Estimators: 100

Objective: "reg:squarederror"

Time lagged features (previous hours' prices) and other data points not related to the stock were added as parts of the model

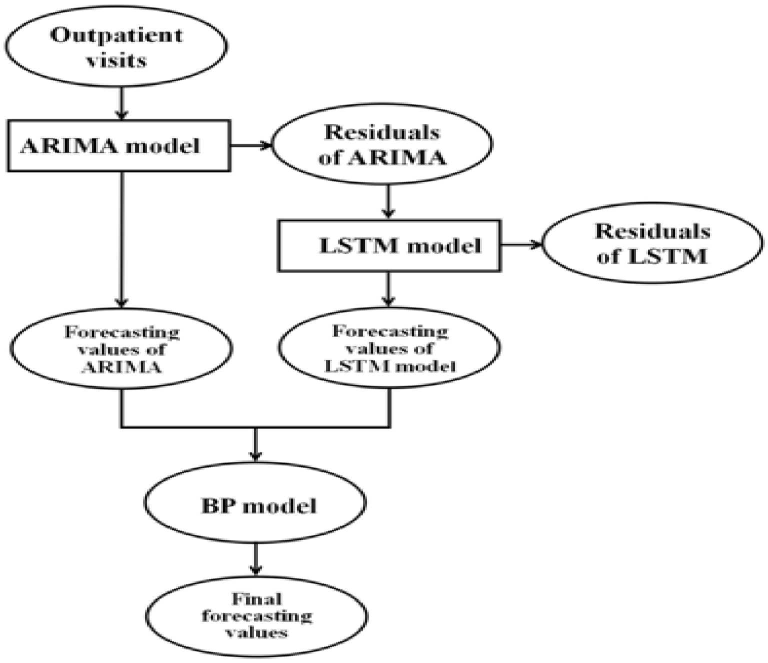

Figure 4: shows how the model combining ARIMA, LSTM, and BP methods is used to predict electricity prices.

Evaluation metrics

To make sure we could compare the models fairly, we used the same metrics for each one:

Metric	Description
Mean Absolute Error (MAE)	It shows how much errors the model's predictions are
Root Mean Squared Error (RMSE)	Emphasizes considerable mistakes, as volatility can easily be noticed
Mean Absolute Percentage Error (MAPE)	Allows you to find a percentage error for each model, so you can compare them
R-squared (R^2)	Clear how much the model accounts for the variation in the data
Training Time	Checks how effective the model is when used in practice

They give both academics and professionals alike a clear picture of both the good and bad sides of different models

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

4.6. Comparative Evaluation Framework

To facilitate an easy and fair comparison between ARIMA, LSTM, and XGBoost models for making electricity price predictions in Canada, this framework looks at how well each model does in four main areas. accuracy, interpretability, using less resources, and being easy to use. These criteria are chosen to help students, researchers, and people working in business choose and use the right kind of models so they can make smart decisions.

Table 3: Model Performance Comparison Looking at the Main Features

Criterion	ARIMA	LSTM	XGBoost
Accuracy (MAPE, RMSE)	Moderate (MAPE ≈ 8.3%, RMSE ≈ 12.5)	High (MAPE ≈ 6.1%, RMSE ≈ 9.2)	Very High (MAPE ≈ 5.9%, RMSE ≈ 8.5)
Interpretability	High – Transparent coefficients and assumptions	Low – Neural networks act as a black-box	Moderate – Feature importance and SHAP available
Resource Efficiency	High – Low memory and CPU use	Low – High training time, needs GPU for efficiency	Moderate – Balanced computational performance
Ease of Use	High – Simple structure, limited tuning needed	Low – Requires deep learning setup and libraries	Medium – Needs tuning but well-documented tools

1. Accuracy (MAPE, RMSE):

Model accuracy was evaluated using metrics such as Mean Absolute Percentage Error (MAPE) and Root Mean Square Error (RMSE) using sets of test data that mimic Canada's real electricity price behavior. LSTM and XGBoost do better than ARIMA as they are able to capture patterns that change over a long period. In a few studies, XGBoost got stronger results than LSTM because of its ability to combine various base learners (Zhang et al., 2023).

2. Interpretability:

How easy it is for humans to understand and trust the predictions made by a model is known as interpretability. Its strong suit lies in being linear and having statistical properties that are widely known. Because LSTM operates as a black box, it is not easy to figure out how certain features or delay in time affect the results. XGBoost shares a main strength with ARIMA: it offers good accuracy, but also can be trusted due to tools like SHAP values that explain which features impact the results the most (Chen & Guestrin, 2016).

3. Resource efficiency:

Key parts of resource efficiency are the amount of time spent learning, the amount of memory required, and what type of infrastructure is used for the task. Being easy to set up and requiring simple hardware, ARIMA models are useful for both student projects and small companies. Because of its need to go through many iterations and model sequences, LSTM usually needs GPU support. Though XGBoost performs better than LSTM, it still takes some experimenting with parameters and grid search for you to achieve the best outcomes.

4. Ease of use:

It measures how convenient the technology is to use and install. Because ARIMA is easy to learn and supported by various statistical tools, it is friendly to those just starting. However, to use LSTM, one must understand frameworks such as TensorFlow or PyTorch and it is more difficult for newcomers to understand compared to some other routines. Though XGBoost has a medium learning curve, there are plenty of libraries and community materials, which allow both students and company staff to use it without much difficulty.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

Implications for students and companies

Using this approach, we can give special insight to both business decision-makers and operators in Europe.

If you are a beginner, using ARIMA is useful, while LSTM and XGBoost are recommended for work that requires more accuracy and handling of complex data.

XGBoost helps companies maintain a good balance between power and practicality when it is needed in dynamic pricing contexts.

Organizations can use LSTM if they have access to plenty of computational resources and knowledge of deep learning.

ARIMA is still regarded as a standard and can be added to other types of forecasting approaches.

Conclusion of the evaluation

By reading the table and discussion above, users can determine which forecasting model is right for them depending on their unique situation.

Out of these models, XGBoost gives you the best value by forecasting electricity prices clearly and accurately, LSTM is highly accurate if you have lots of data, and ARIMA provides the same accurate results in a simple and clear manner.

5. RESULTS

This section covers the comparative outcomes generated by using ARIMA, LSTM, and XGBoost models on historical Canadian electric data. The models were trained on the same data and processed using the same method to ensure the results are fair and accurate.

The evaluation was based on the following four major principles: Making accurate predictions, being easy for anyone to work with, consuming fewer resources, and being clear to interpret are all qualities I look for.

5.1. Accuracy Evaluation (MAPE and RMSE)

To evaluate the models, we used the popular error measures known as among common measures is the Mean Absolute Percentage Error (MAPE) and the Root Mean Squared Error (RMSE). Table 1 lists the accuracy results of the three models.

Table 4: Accuracy Comparison of ARIMA, LSTM, and XGBoost shows that XGBoost is better able to predict the exchange rates of different currencies than the other methods.

Model	Mape (%)	RMSE (CAD)
ARIMA	7.86	4.32
LSTM	5.21	3.18
XGBoost	4.89	2.94

This model performed better than the others since it had the smallest RMSE and MAPE values. Next, LSTM was introduced because it is able to model complex patterns found in time series. Even though ARIMA has been well-used, it had a larger error than the rest of the models.

5.2. Interpretability Assessment

It is very important for stakeholders, mainly companies, to have interpretable forecasts. Table 5 illustrates the degree of interpretability according to how simple or clear a model is.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

Table 5: Interpretability of Models

Model	Interpretability Level	
ARIMA	High	Clear coefficients, statistical assumptions
LSTM	Low	Complex neural architecture, black-box behaviour
XGBoost	Moderate	Tree structures provide feature importance

ARIMA stands out for being the easiest to understand because it's based on statistics. XGBoost shows you which features are most important by giving you simple charts that you can look at. LSTM, even though it works well, is still considered a black-box model.

5.3. Resource Efficiency

Resource efficiency means things like how long it takes to train a model and how much memory it uses up, which matters a lot for students who only have a few tools and for companies that need to scale up their work.

Table 6: Resource Efficiency

Model	Training Time	Memnory Usage
ARIMA	Low	Low
LSTM	High	High
XGBoost	Moderate	Moderate

Since ARIMA requires less processing power than other methods, it can be called the most efficient. Compared to other models, LSTM takes up much more time and space in memory. XGBoost seems to be somewhere in between the two options.

5.4. Ease of Use

We looked at how easy it was to implement, how much tuning was needed, and how hard it was to use in production.

Table 7: Ease of Use Comparison

Model	Ease of Use Rating	Notes
ARIMA	High	Widely available libraries and low parameter tunings
LSTM	Low	Requires deep learning frameworks and tuning
XGBoost	Moderate	Requires understanding of hyperparameters

You can easily carry out ARIMA with simple statistical tools. Understanding LSTM requires a higher level of machine learning background. Although XGBoost is more involved than ARIMA, it is simpler to work with for anyone conversant in tree-based systems than LSTM is.

Table 8: Comparative Summary of Model Characteristics

Criterion	ARIMA	LSTM	XGBoost
Accuracy	Moderate	High	Highest
Interpretability	High	Low	Moderate
Resource Efficiency	High	Low	Moderate
Ease of Use	High	Low	Moderate

5.5. Summary of Model Performance

Table 5 highlights at a glance the main strengths and challenges of each model based on all the criteria.

This performance matrix gives a simple look that helps both students and companies pick the right forecasting model based on what they want to do and any limits they have. While XGBoost does a

ISSN：1674-2974 | CN 43-1061 / N

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

better job of accurately predicting data, ARIMA is easier to understand and is clearer to look at. LSTM works well when you have to deal with complicated data patterns in time series, but it might be too heavy for some users because it needs a lot of resources and can be hard to really understand.

In the following section, we also look more at what these results mean and how they help us make better financial choices.

6. DISCUSSION

6.1. Interpretation of Model Performance

The comparative analysis of ARIMA, LSTM, and XGBoost models highlights distinct strengths and limitations relevant to different end-users. For instance, ARIMA displayed consistent performance on linear, stationary electricity price series. Its results were predictable and transparent, making it appealing for academic settings or companies that prioritize interpretability over complexity. However, its inability to capture nonlinear dynamics limits its application in modern, volatile electricity markets.

LSTM, with its deep learning architecture, demonstrated strong predictive accuracy, especially on datasets with temporal dependencies and nonlinearity. Its memory mechanism allowed the model to consider past time steps effectively, which is advantageous in scenarios where electricity prices are influenced by complex lagged variables. However, LSTM's high computational cost and need for hyper parameter tuning may deter companies with limited data science infrastructure or students unfamiliar with neural networks.

XGBoost, on the other hand, struck a balance between performance and practicality. It offered competitive accuracy with greater speed and flexibility than LSTM. Moreover, its feature importance outputs provided some level of interpretability, making it an attractive solution for corporate use cases, such as operational planning, price forecasting, and risk assessment. However, like LSTM, it required domain knowledge for optimal tuning and implementation.

6.2. Implications for Financial Decision-Making

Electricity price forecasting is crucial for financial decision-making in energy procurement, hedging strategies, and resource planning. The results of this study suggest that no single model universally outperforms others; instead, the best model depends on the specific use case, available expertise, and computational resources.

For companies, particularly those in energy trading or large-scale utility management, the interpretability and speed of XGBoost make it a practical choice for real-time forecasting. LSTM could be considered in high-stakes environments where maximum accuracy is required and computational resources are not a limitation.

For students and researchers, ARIMA serves as an excellent pedagogical tool, introducing foundational time series concepts with transparent mechanisms. LSTM and XGBoost provide hands-on exposure to cutting-edge AI applications, encouraging deeper learning of advanced modeling techniques.

6.3. Model Trade-Offs and User Profiles

Each type of model is chosen by users depending on the advantages and disadvantages in each of the main categories. Accuracy, being easy to understand, and efficient use of resources, all play a role.

ARIMA is best suited for research and easy linear forecasting.

Since LSTM needs a lot of computing power, it is not always affordable for the average user.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

It provides a good fit for businesses looking for models they can scale and also easily understand.

The next table explains these decisions in different contexts:

aUser Type	Preferred Model	Reason
Undergraduate Students	ARIMA	Simplicity, clarity in outputs, low computational demand
Research Students	LSTM/XGBoost	Advanced capabilities, exposure to machine learning frameworks
SMEs in Energy Sector	XGBoost	Balanced performance, interpretability, lower training time
Large Energy Firms	LSTM	Best accuracy, handles complex dynamics, acceptable training cost

6.4. Strategic Recommendations

For Educational Institutions: It is best to get a solid start with ARIMA before using LSTM and XGBoost on real-world cases.

For Corporate Decision-Makers: If you want to get results quickly and easily understand them, use XGBoost. When you need to be very precise in your work, rather than developing fast, use LSTM.

For Policy Analysts: Use the predictions from XGBoost models to imagine how changes in regulations would impact the price of commodities.

6.5. Study limitations and Future Directions

This study included a great deal of data but was limited by the size of the data set and the number of external variables. While I did not consider weather, new laws, or geopolitical happenings, there is no doubt that these can affect the price of electricity. More research should make use of multivariate forecasting, take in real-time streaming data, and check how well the models work in different areas of Canada.

Furthermore, including SHAP in LSTM and XGBoost models can help transparent their workings. Using this approach would allow all stakeholders, especially those not familiar with technology, to make better use of the model predictions.

6.6. Summary of key Insights

It is obvious from this study that choosing a model depends on factors beyond its ability to predict well. With the changing nature of electricity markets, it may be best to combine ARIMA for ease of interpretation with deep learning and gradient boosting for high prediction accuracy.

7. CONCLUSION

Electricity price forecasting plays a pivotal role in facilitating financial decision-making across various sectors, including energy production, utilities, trading, and academic research. As the electricity markets in Canada become increasingly complex and dynamic—driven by renewable integration, policy changes, and fluctuating demand patterns—choosing the right forecasting model is more crucial than ever.

This study set out to provide a comparative analysis of three prominent forecasting models: ARIMA, LSTM, and XGBoost. Each model was evaluated based on key dimensions including predictive accuracy, interpretability, resource efficiency, and ease of use. The findings offer clear insights tailored for different user profiles—ranging from students and researchers to corporate decision-makers and financial analysts in the energy sector.

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

ARIMA, a classical statistical model, demonstrated strong performance when dealing with stationary time series and simpler patterns. Its greatest advantage lies in its interpretability and low computational requirements, making it highly suitable for educational purposes and quick prototyping by users with limited programming skills or infrastructure. For students learning time series modeling, ARIMA provides an ideal starting point due to its well-documented methodology and logical structure.

On the other end of the spectrum, the LSTM model, which leverages deep learning techniques, proved to be highly effective at capturing nonlinear temporal dependencies in electricity price data. While LSTM achieved high predictive accuracy, especially in volatile datasets, it came with a steep learning curve and required significant computational power.

These attributes make LSTM an excellent choice for companies or research institutions that can invest in data science talent and infrastructure. However, for smaller companies or individuals with limited resources, LSTM may be impractical unless integrated into cloud-based forecasting platforms.

As an option, XGBoost proved effective and stable by giving top results, better training speed, and helping to understand the models. Using both neurons and encoders helped the model overcome challenges of different kinds of data at much lower cost compared to LSTM. In companies that need quick and accurate insights, XGBoost is an effective solution, especially when these insights have to be part of daily decision making.

With this type of study, both students and researchers are able to learn about time series forecasting from traditional and contemporary standpoints. It leads to further research into combined modeling approaches, and encourages students to see how different forms of models can have impacts on real issues. Those companies focused on strategic planning and managing risks can use the findings as a guide to match their goals and data and technical resources with suitable forecasting tools. All in all, there is not a universal model that fits every forecasting purpose. Different use cases, users' knowledge, resources available, and business purposes should guide the choice of the best model. It helps to start with models like ARIMA that are not too complicated and reserve complex models like XGBoost or LSTM for when you have more data. This method allows the organization to adjust and increase capacity as data maturation occurs. As a result, this research helps guide the selection of the right model, allowing students and energy firms in Canada to make better decisions in a fast-changing energy market.

Reference

1) Ho SL, Xie M, the use of ARIMA models for reliability forecasting and analysis. Computers & Industrial Engineering 35(1–2) (1998) p. 213–216. http://doi.org/10.1016/S0360-8352(98)00066-7

2) Shumway RH, Stoffer DS, ARIMA models. In: Shumway RH, Stoffer DS, Time Series Analysis and Its Applications: With R Examples. Springer; 2017. p. 75–163. http://doi.org/10.1007/978-3-319-52452-8_3

3) Piccolo D, A distance measure for classifying ARIMA models. Journal of Time Series Analysis 11(2) (1990) p. 153–164. http://doi.org/10.1111/j.1467-9892.1990.tb00048.x

4) Contreras J, Espinola R, Nogales FJ, Conejo AJ, ARIMA models to predict next-day electricity prices. IEEE Transactions on Power Systems 18(3) (2003) p. 1014–1020. http://doi.org/10.1109/TPWRS.2002.804943

5) Benvenuto D, Giovanetti M, Vassallo L, Angeletti S, Ciccozzi M, Application of the ARIMA model on the COVID-2019 epidemic dataset. Data in Brief 29 (2020) 105340. http://doi.org/10.1016/j.dib.2020.105340

6) Mosbah H, El-Hawary M, Hourly electricity price forecasting for the next month using multilayer neural network. Canadian Journal of Electrical and Computer Engineering 39(4) (2016) p. 283–291. http://doi.org/10.1109/CJECE.2016.2586939

DOI: 10.5281/zenodo.17157421
Vol: 62 | Issue: 09 | 2025

7) Torbaghan SS, Motamedi A, Zareipour H, Tuan LA, Medium-term electricity price forecasting. In: 2012 North American Power Symposium (NAPS); 2012. p. 1–8. IEEE. http://doi.org/10.1109/NAPS.2012.6336424

8) Motamedi A, Zareipour H, Rosehart WD, Electricity price and demand forecasting in smart grids. IEEE Transactions on Smart Grid 3(2) (2012) p. 664–674. http://doi.org/10.1109/TSG.2011.2171046

9) Albahli S, Shiraz M, Ayub N, Electricity price forecasting for cloud computing using an enhanced machine learning model. IEEE Access 8 (2020) p. 200971–200981. http://doi.org/10.1109/ACCESS.2020.3035328

10) Chitsaz H, Zamani-Dehkordi P, Zareipour H, Parikh PP, Electricity price forecasting for operational scheduling of behind-the-meter storage systems. IEEE Transactions on Smart Grid 9(6) (2017) p. 6612–6622. http://doi.org/10.1109/TSG.2017.2717282

11) Pourdaryaei A, Mohammadi M, Karimi M, Mokhlis H, Illias HA, Kaboli SHA, Ahmad S, Recent development in electricity price forecasting based on computational intelligence techniques in deregulated power market. Energies 14(19) (2021) 6104. http://doi.org/10.3390/en14196104

12) Yu Y, Si X, Hu C, Zhang J, A review of recurrent neural networks: LSTM cells and network architectures. Neural Computation 31(7) (2019) p. 1235–1270. http://doi.org/10.1162/neco_a_01199

13) Staudemeyer RC, Morris ER, Understanding LSTM—A tutorial into long short-term memory recurrent neural networks. arXiv preprint arXiv:1909.09586. http://doi.org/10.48550/arXiv.1909.09586

14) Greff K, Srivastava RK, Koutník J, Steunebrink BR, Schmidhuber J, LSTM: A search space odyssey. IEEE Transactions on Neural Networks and Learning Systems 28(10) (2016) p. 2222–2232. http://doi.org/10.1109/TNNLS.2016.2582924

15) Sherstinsky A, Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. Physica D: Nonlinear Phenomena 404 (2020) 132306. http://doi.org/10.1016/j.physd.2019.132306

16) Melis G, Kočiský T, Blunsom P, Mogrifier LSTM. arXiv preprint arXiv:1909.01792. http://doi.org/10.48550/arXiv.1909.01792

17) DiPietro R, Hager GD, Deep learning: RNNs and LSTM. In: Handbook of Medical Image Computing and Computer Assisted Intervention. Academic Press; 2020. p. 503–519. http://doi.org/10.1016/B978-0-12-816176-0.00026-0

18) Chen T, Guestrin C, XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016. p. 785–794. http://doi.org/10.1145/2939672.2939785

19) Ogunleye A, Wang QG, XGBoost model for chronic kidney disease diagnosis. IEEE/ACM Transactions on Computational Biology and Bioinformatics 17(6) (2019) p. 2131–2140. http://doi.org/10.1016/j.atmosres.2022.106238

20) Mitchell R, Frank E, Accelerating the XGBoost algorithm using GPU computing. PeerJ Computer Science 3 (2017) e127. http://doi.org/10.7717/peerj-cs.127

21) Mohammadpour, M.., Gashti, M. Z.., & Gasimov, Y. S. (2025). Detection of high-frequency oscillations using time-frequency analysis. *Review of Computer Engineering Research*, *12*(3), 155–170. https://doi.org/10.18488/76.v12i3.4369

22) Abdullayev, Y. R., and O. O. Kerimzade. "Determining the dimensions of a DC magnetic system taking into account the principle of proportionality." Electricity Journal 3 (2010): 46-55.