Journal of Statistics Applications & Probability
An International Journal

http://dx.doi.org/10.18576/jsap/150330

A Comparative Statistical Study of Time Series and Deep Learning Models for Short-Term Stock Price Forecasting

Sarah Aljohani[1], Amir Ahmad Dar[2,] and I. Silambarasan[3]*

[1] Department of Mathematics and Sciences, Prince Sultan University, Riyadh, 11586, Saudi Arabia
[2] Department of Statistics, Lovely Professional University, Punjab, India
[3] Department of Mathematics, Vel Tech Multi Tech Dr. Rangarajan Dr. Sakunthala Engineering College, Avadi, Chennai, Tamil Nadu, India

Received: 23 Feb. 2026, Revised: 2 Apr. 2026, Accepted: 22 pr. 2026.
Published online: 1 May 2026.

Abstract: In this study, we investigate the performance of different predictive models in forecasting the closing stock price of RTPower Ltd. over a 15-day horizon, a task with significant implications for investment policy and risk management given the centrality of financial prediction to modelling practice. Historical data on daily stock prices from the National Stock Exchange was gathered and pre-processed, and training and testing of different models were carried out, including time-series and machine learning algorithms. The models were evaluated based on standard evaluation metrics such as MAE, MSE, RMSE, and MAPE. It was shown that LSTM was more successful in comparison with other models because it showed the lowest percentage of errors and the highest degree of predictive success. LSTM achieved the lowest errors across all metrics (MAE: 1.4621, RMSE: 1.7458, MAPE: 11.17%), outperforming SARIMA (MAPE: 25.64%), XGBoost (MAPE: 14.71%), Prophet (MAPE: 38.37%), and Transformer (MAPE: 43.27%). This is possibly due to its capability to learn temporal dependencies in the data. The chosen model was applied to determine the closing prices for the following 15 days, and it proved highly practical for short-term trading strategies. This paper also highlights the significance of higher-level modelling tools in predicting financial time series and, simultaneously, underscores the necessity of selecting a model suited to the specific prediction problem. The constraints involving susceptibility to market fluctuations and data quality are addressed, along with suggestions for future studies, including the incorporation of external variables such as market sentiment. The findings contribute to the existing literature on stock price prediction and offer practical guidance to investors.

Keywords: Predictive Modelling, LSTM, ARIMA, SARIMA, XGBoost, Transformer.

1. Introduction

The stock price prediction is one of the foundations of the financial analysis on which investors, traders, and portfolio managers can build the appropriate strategies to follow through uncertainties in the market [1]. This ability to predict the closing price with a good level of precision especially on small frames like 15 days is very important in the formulation of trading strategies, optimization of assets, and the risk management [2]. This has gained more and more relevance in modern markets where for full and swirling price changes that occur as a result of economic events, investor sentiment and global factors and influences makes precise and timely forecasts extremely attractive [3]. Nevertheless, the inherent complexity of stock prices dynamics that imply nonlinearity, seasonality, and stochastic dynamics renders conventional forecasting extremely difficult [4]. To illustrate, an ARIMA model, which is often regarded as one of the most significant tools of time-series analysis, does not often manage to capture seasonality trends, or even abrupt shifts in the financial data, therefore restricting its applicability to dynamical market conditions [5-6]. The choice of the research motivation is based on a wish to contribute to bridging the gap between the new theoretical achievements in predictive modeling and their application on financial markets. Empirical evidence of the effectiveness of models in seasonal time series is provided by using a realistic dataset and a wide range of models based on the ongoing debate on the topic of financial forecasting. In the light of this, the findings are supposed to make efficient guidelines to the investors and analysts who require good forecasts that would help them in the short-term and relative evaluation of the advantages and disadvantages of each method. This paper is organized in the way it presents a literature review, followed by the methodology, findings, and discussion with references to the way the research and practice impacts the future.

The main area of research in financial analytics is stock price forecasting because it has a direct effect on both risk management and investment decisions. Researchers like the ARIMA and SARIMA models because they are traditional time

*Corresponding author e-mail: amirphd111@gmail.com

series forecasting methods that have strong theoretical support and are easy to understand. Box and Jenkins created ARIMA as a standard tool for predicting linear time series. But a number of real-world studies have shown that these models have problems with financial markets because they need to be able to handle both nonlinear movements and sudden changes in the market [6-7].

Ensemble models like SVM and XGBoost are now important machine learning methods that help to solve these problems. Chen and Guestrin (2016) showed that XGBoost works well with structured financial data and can handle large amounts of it [8]. Patel et al. (2015) showed that hybrid models that combine technical indicators with machine learning algorithms do better than traditional statistical methods when it comes to predicting stocks. Facebook's Prophet is a tool that lets users do seasonal and trend analysis on financial time series data. It is easy to use and has a lot of options [9].

Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, which Sepp Hochreiter and Jürgen Schmidhuber made, have helped the field of deep learning get better results. LSTM models are great at predicting stock prices because they can keep track of long periods of time and look at how patterns change over time. Fischer and Krauss (2018) [10] and Li et al. (2023) [11] performed studies showing that LSTM is better than ARIMA and other machine learning models in terms of both accuracy and reliability. Vaswani et al. (2017) [12] work on Transformer-based architectures has led to the creation of attention mechanisms that improve the modeling of financial datasets by being able to handle complex patterns that change over time.

The increasing importance of comparative research aids in identifying the most effective models for diverse market conditions. Lim and Zohren (2021) [13] and Khadka et al. (2021) [14] show that no one model works better than all the others in every test situation. This is because the results of a model depend on the data's characteristics, the methods used to preprocess it, and the time frame for making predictions. Researchers are now concentrating on hybrid and ensemble methods, as these approaches leverage the strengths of both statistical and deep learning models to improve forecasting accuracy.

The researchers encounter three significant challenges in their work due to the necessity for their system to manage data noise, market fluctuations, and external influences, including economic indicators and the attitudes of market participants. The researchers discovered that integrating news sentiment with social media signals as supplementary data sources enhanced their predictive capabilities regarding forthcoming events. The current research indicates that statistical models have evolved into more sophisticated machine learning and deep learning techniques [15-16]. The study found that LSTM and Transformer-based models do a better job of predicting short-term stock price movements than ARIMA and SARIMA benchmarks. This is because they can better capture complex relationships.

The rise of deep learning has opened up powerful new possibilities for forecasting financial markets. Studies comparing LSTM, GRU, and Transformer models on Tesla stock data spanning 2015 to 2024 have found that LSTM consistently stands out, achieving around 94% accuracy while reliably capturing underlying market trends. Researchers have also explored more sophisticated ensemble approaches that combine Variational Autoencoders, Transformers, and LSTM networks together, with the idea that each model brings something different to the table collectively handling both linear patterns and more complex, non-linear price behaviors across a range of datasets. Perhaps most interestingly, recent work has shown that blending LSTM's ability to track price movements with Transformer-based sentiment analysis through tools like FinBERT can outperform models that rely on price data alone, pointing to how much market mood and narrative can influence stock behavior. Taken together, these findings reflect just how quickly deep learning methods are reshaping financial forecasting, and they provide a strong motivation for the kind of comparative analysis carried out in this study [17-18].

2. Materials and Methods

This research under consideration is aimed at providing a comparison of the performances of some predictive models to estimate the daily closing prices of a stock at RTPower Ltd. over a period of 15 days. RTPower Ltd. was chosen for this study because it represents an interesting case in the renewable energy space, a company that is growing quickly, publicly listed on the National Stock Exchange, and whose stock tends to move in ways that are anything but predictable. That inherent volatility is precisely what makes it a compelling choice, as it puts the forecasting models to a real test and offers a clearer picture of how well each one holds up when market conditions are constantly shifting. These steps include the process of data collection, data cleaning and preprocessing, EDA, model selection, model training, and model performance evaluation to be able to make systematic identification of the optimal forecasting method. These steps have been summarized in a flowchart as follows so as to provide visual representation of the process as illustrated in Figure 1.

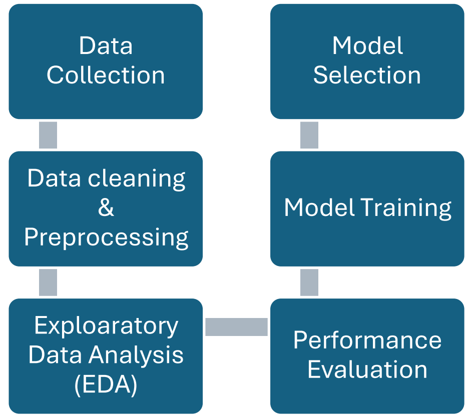

Fig. 1: Methodology flowchart illustrating the end-to-end research pipeline

2.1 Data Collection

A dataset contains daily stock price records from equity stocks which cover 733 observations between March 29 2022 and March 17 2025. The data set includes six variables which are date, opening price, high, low, previous close, last traded price (LTP) and closing price, volume-weighted average price (VWAP), 52-week high and low, volume, value and the number of trades. The researchers selected close as their main prediction target because it serves as an essential metric for financial assessments. The information was obtained at NSE. The models developed from this data will maintain their strength through all market situations which include both volatility and seasonal market patterns and thus prove their operational effectiveness.

2.2 Data Cleaning and Preprocessing

Once the datetime formatting was in place, the closing price column was checked to make sure there were no unexpected gaps in the data. On the few occasions where trading days were missing, the last known value was carried forward, which helped maintain the natural flow of the time series without distorting it. A check for duplicate records came up clean. To get a sense of whether any unusual values needed to be dealt with, the data was visually examined through a time series line plot and a rolling standard deviation chart. While the 2022–2023 period showed some noticeable swings, these reflected real market behavior rather than errors, so no data points were removed.

Stationarity was then formally tested using the Augmented Dickey-Fuller test, which confirmed that the original series was non-stationary. To address this, first-order differencing was applied, and the series passed the stationarity check afterward — a necessary step before fitting any ARIMA-family models. For the LSTM and XGBoost models, however, the raw closing price series was kept intact as the target variable, with Min-Max normalization used to bring values into the range of 0 to 1. Importantly, the scaler was fitted only on the training data and then applied to the test set separately, ensuring that no information from the test period crept into the preprocessing stage and potentially skewed the results.

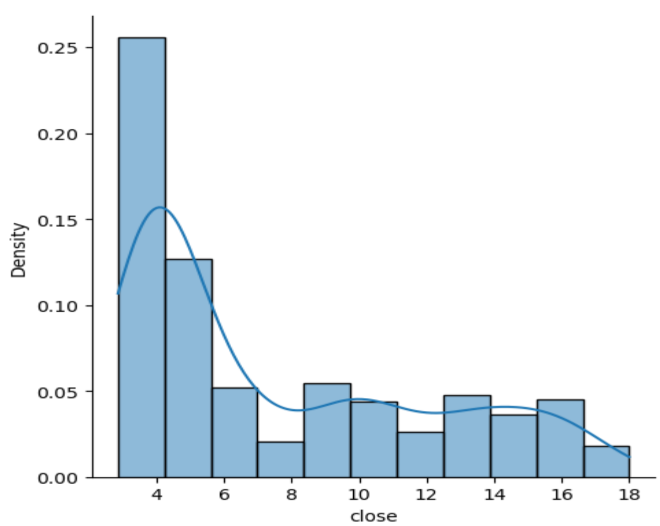

Fig. 2: Distribution plot of the daily closing price of RTPower Ltd.

2.3 Exploratory Data Analysis (EDA)

The researchers conducted extensive EDA procedures to establish a complete comprehension of the primary structure and features which compose the closing price data. The distribution of the closing price presented in figure 2 allows viewers to see central tendency and variability and skewness through its various distribution patterns.

Besides, a time series line plot of the closing prices was drawn, which reflected the trends and probable seasonal patterns in Figure 3.

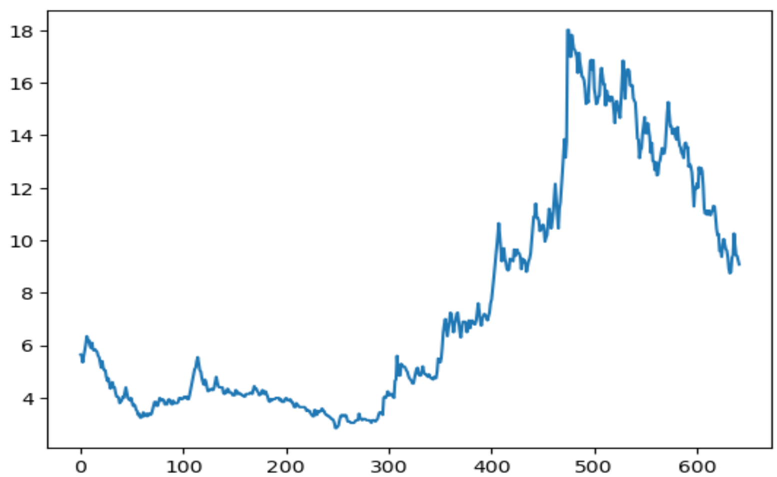

Fig. 3: Time series line plot of the daily closing prices of RTPower Ltd. over the full study period

To further investigate the temporal nature of the series, the rolling statistics were also computed and shown, the rolling mean and rolling standard deviation. In this regard, based on figure 4, it can be stated that the rolling metrics revealed that the mean and variance was not constant throughout the time series, thereby showing non-stationarity in the time series.

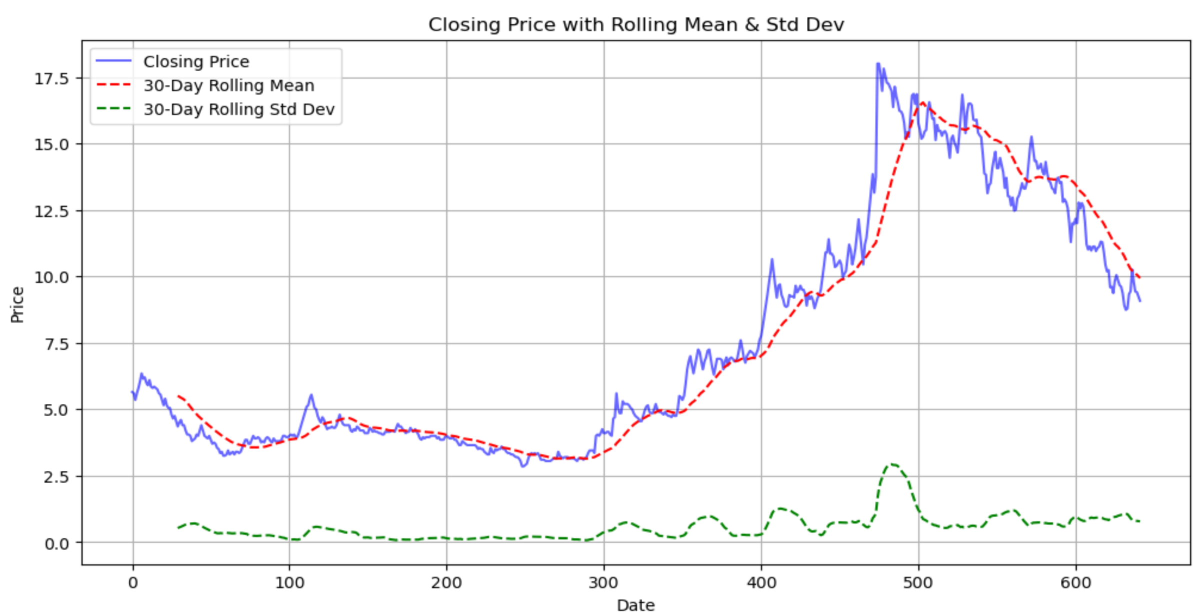

Fig. 4: Closing price of RTPower Ltd. plotted alongside its 30-day rolling mean and rolling standard deviation, demonstrating non-stationarity in the series as evidenced by the varying mean and fluctuating variance across the observation period

This experimental finding was to be established further on with formal stationarity tests that verified that the original series was non-stationary. The first-order differencing was then applied to overcome this problem and this converted the series into stationary scale, which could be analyzed using time series. In addition, the autocorrelation and partial autocorrelation function plots were plotted to identify potential lag dependencies as well as inform the selection of the model as shown by Figure 5.

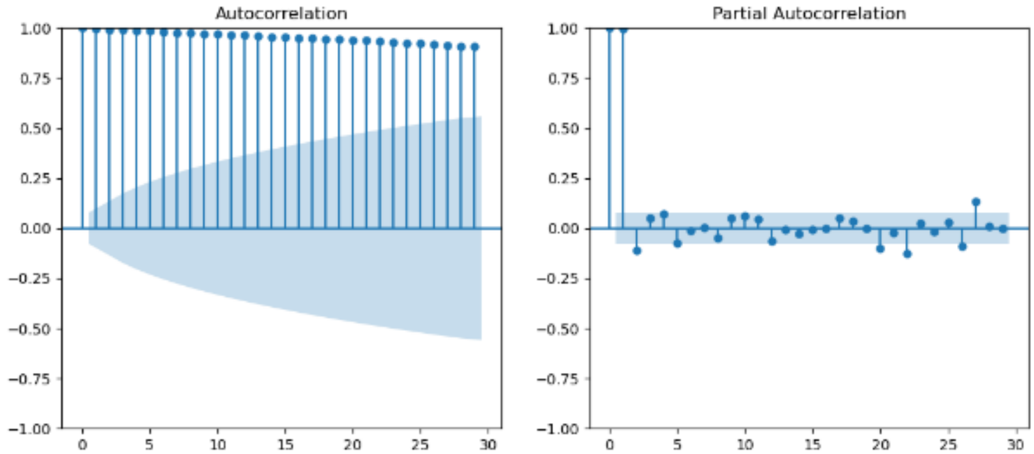

Fig. 5: ACF and PACF plots of the differenced closing price series

Additionally, a time series seasonal decomposition was conducted, which clearly separated the trend, seasonal, and residual parts of the data as illustrated in Figure 6.

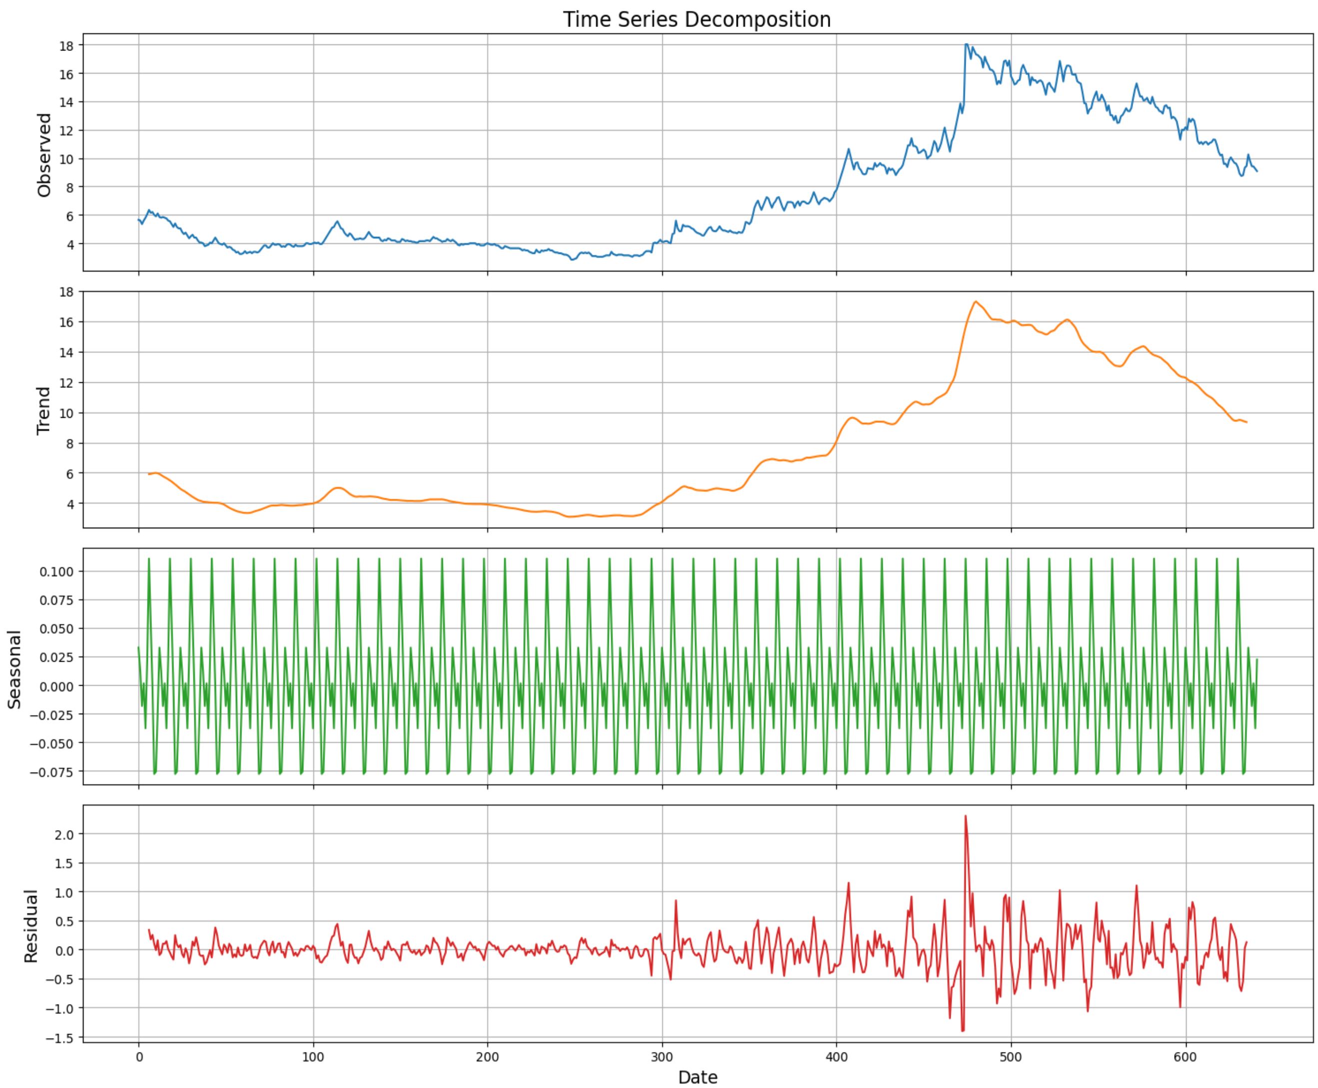

Fig. 6: Seasonal decomposition of the closing price time series into its three constituent components

All these analyses thereby provided a sound basis for selecting appropriate forecasting models and ensured that key assumptions such as stationarity were properly met before embarking on actual model development.

2.4 Model Selection

Various models of forecasting were selected in this study to evaluate and compare their ability in predicting stock closing price. This was to integrate classical statistical models and the latest machine learning and deep learning algorithms in determining the most effective mechanism of capturing the temporal dynamics and underlying data patterns. First of all, six models have been considered: ARIMA [19], SARIMA [20], LSTM [21], Prophet [22], XGBoost [19, 23] and Transformer.

The five models were chosen deliberately to cover a wide range of approaches to financial time series forecasting, rather than relying on any single methodology. SARIMA was included as a classical baseline, it has a long track record of handling autocorrelation and seasonal patterns in linear data, making it a natural starting point for comparison [24]. Prophet was brought in for a different reason: its ability to cleanly separate trend and seasonality components without demanding heavy preprocessing makes it particularly well-suited to financial data that does not stay still [24]. XGBoost earned its place through its consistent performance on structured financial data, especially when paired with carefully engineered lag features and rolling statistics that give the model a richer sense of recent price behavior [8]. LSTM was chosen as the core deep learning model because of how its gating mechanisms allow it to hold onto relevant information over longer sequences, a quality that matters a great deal when trying to capture the shifting, non-linear dynamics of stock prices [25]. Finally, the Transformer model was included to see how well self-attention mechanisms perform on this particular dataset, given the growing body of evidence suggesting their strength in forecasting across multiple time horizons [12, 18].

2.5 Model Training

The forecasting methods that were selected through the research process which included SARIMA, LSTM, Prophet XGBoost and Transformer were subjected to independent training based on historical closing price data after the model selection process was completed. The data was split into training and testing sets in chronological order rather than shuffled randomly which is simply the right way to handle time series data, since the whole point is to predict what comes next, not what came before. Keeping the split sequential ensures the model only ever learns from the past and is tested on genuinely unseen future data. Equally important is the principle that any preprocessing transformations, such as scaling, must be learned exclusively from the training set. Allowing test data to participate in any fitting step, even indirectly, lets future information seep into the model's learning process, which quietly undermines the reliability of the evaluation results. The SARIMA model with the selection of the correct parameters p d q and seasonal parameters P D Q s was developed through the ACF and PACF plot analysis together with seasonal decomposition method analysis. AIC minimization with residual analysis was used to identify the optimal model which would enhance model simplicity while maintaining adequate performance. The LSTM model required data inputs to be standardized through Min-Max scaling because this model depends on input size and scaled data would boost its learning capability. The researchers developed a supervised learning method from time series data by implementing a rolling window method. The system architecture includes stacked LSTM layers which use dropout regularization to prevent overfitting together with dense layers that link outputs to target predictions. The training process involved multiple epochs which used the Adam optimizer together with MSE loss function while early stopping was implemented to end training when the model achieved convergence. The Prophet model needed minimal preprocessing before it directly processed closing price data which included additive components that helped identify trends and seasonal patterns. The model successfully forecasted results because it contained built-in methods for handling holidays and changepoints and it managed to cope with non-linear and unpredictable time sequences. The XGBoost model achieved successful operation through its creation of lag features combined with rolling statistics which produced different performance results.

The LSTM model was built with two stacked layers of 64 units each, with a dropout rate of 0.2 applied after every layer to reduce the risk of overfitting, and a single Dense layer at the output end. These choices were not arbitrary, they reflect what the literature has consistently pointed to as a sensible balance for financial forecasting tasks. Research has shown that keeping the number of units per layer moderate, around 64, paired with appropriate dropout, tends to produce reliable predictions without overcomplicating the architecture. At the same time, studies using slightly larger configurations, such as two layers of 128 neurons with a dropout rate of 0.4, have also demonstrated strong results, reinforcing the value of this general design pattern.

The Adam optimizer was selected to handle the training process, largely because of its ability to adapt the learning rate as training progresses, which generally leads to faster and more stable convergence. An initial learning rate of 0.001 was used, consistent with what prior work in financial forecasting has recommended [25]. To give the model enough context to learn from, a rolling window of 60 trading days was used as the look-back period — roughly three months of market activity, which is a well-established choice for short-term forecasting.

Getting the hyperparameters right matters enormously with LSTM models. The number of epochs, batch size, and dropout rate all play a significant role in determining whether the model generalises well or simply memorises the training data. To guard against the latter, early stopping was applied with a patience of 10 epochs, meaning training would automatically halt if the validation loss stopped improving. Table 1 brings together the final configuration of all five models in one place, making it straightforward for others to reproduce the experiments.

Table 1: Summary of Model Configurations and Key Hyperparameters

Model	Key Parameters	Tuning Method	Input Features
SARIMA	p=1, d=1, q=1; P=1, D=1, Q=1, s=52	AIC minimisation + ACF/PACF	Differenced closing price
LSTM	2 layers, 64 units, dropout=0.2, Adam (lr=0.001), epochs=100, window=60	Early stopping (patience=10)	Min-Max scaled closing price
Prophet	Additive mode, yearly seasonality, default changepoints	Automated	Raw closing price + date index
XGBoost	n_estimators=100, max_depth=5, lr=0.05	Validation RMSE minimisation	Lag features (1–10 days), rolling statistics
Transformer	2 encoder layers, 4 heads, FFN=128, dropout=0.1, Adam	Validation loss	Min-Max scaled closing price, window=60

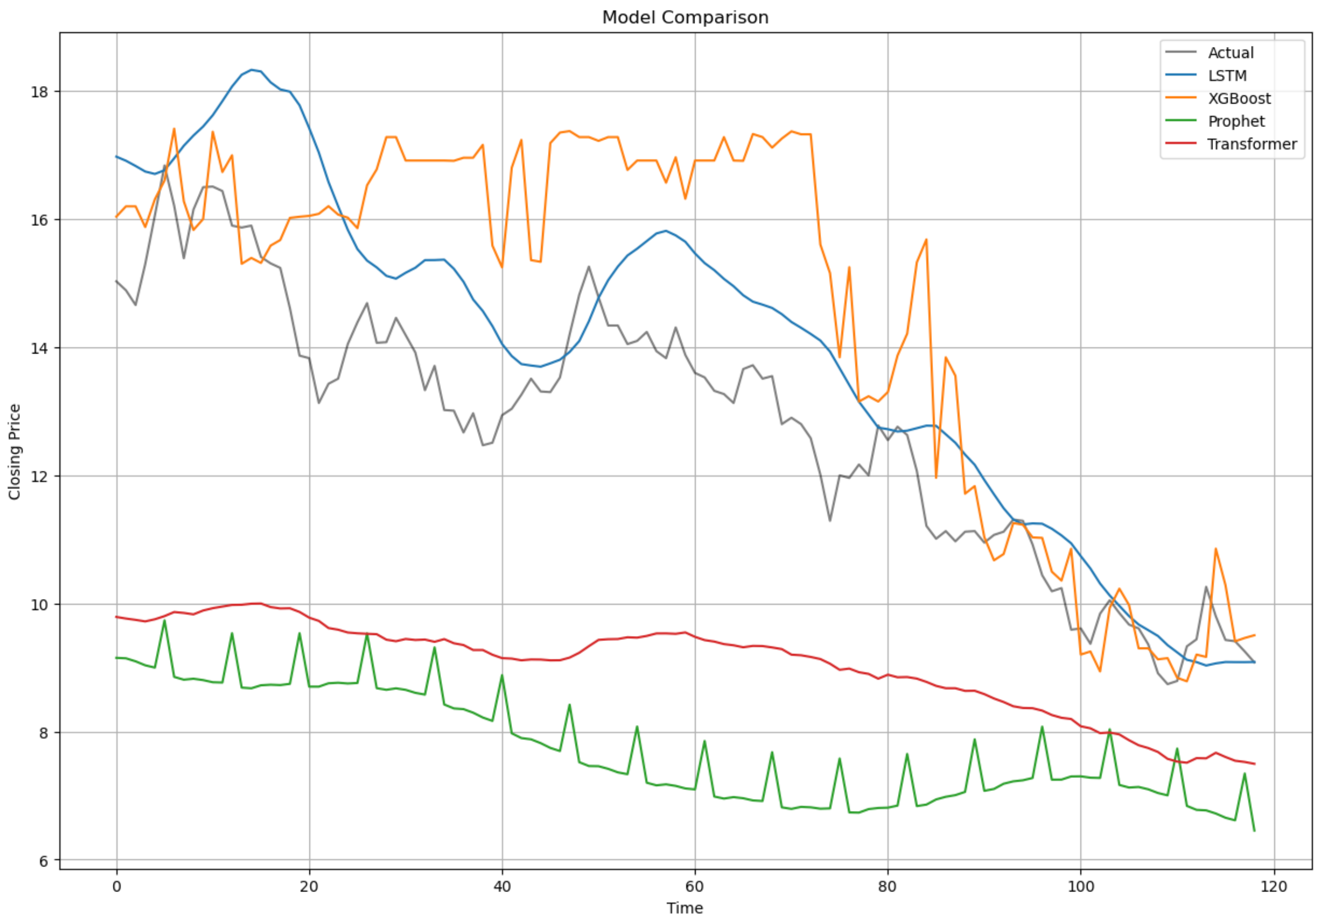

Fig. 7: Comparison of actual vs. predicted closing prices on the test set for all five models (SARIMA, LSTM, Prophet, XGBoost, and Transformer)

2.6 Performance Evaluation

The effectiveness of the selected models is gauged on the basis of four proposed metrics of forecasting, namely, MAE, MSE, RMSE, and MAPE. The four metrics together provide a complete understanding of prediction accuracy and error sensitivity and system resilience. MAE measures the average error size while MSE shows how large errors affect results and RMSE allows users to compare errors with actual data while MAPE shows error magnitude through actual value comparison without any scale dependencies. Among the tested models LSTM model delivered superior performance because it achieved the lowest scores across all evaluation metrics. This model proved effective for modeling stock price movement patterns because it learned to track long-term dependencies while modeling complex nonlinear relationships. The Transformer model demonstrated strong performance because it excelled at detecting intricate time-based relationships yet its results showed less accuracy when compared to LSTM. XGBoost and Prophet demonstrated equal predictive capabilities in their forecasting abilities. Prophet used decomposition methods to create accurate predictions while XGBoost showed effective performance in predicting outcomes during brief time intervals. Although SARIMA works effectively with seasonal data it generates higher error rates which makes it challenging to model the nonlinear and unpredictable nature of financial time series data. The model analysis results show LSTM functions as the most accurate and dependable model for forecasting stock price changes that occur within short time frames when compared to all other models tested in this research. Figure 7 displays the actual prices which are plotted as a line graph together with the prices that the models predicted.

3. Results and Conclusion

Based on the comparison study, one has found out that the LSTM model is the most accurate and reliable model in terms of predicting the stock closing prices. The testing process for five different models, which included SARIMA and LSTM and Prophet and XGBoost and Transformer, employed test parameters that included MAE and MSE and RMSE and MAPE. The LSTM model produced the lowest error rate according to all testing parameters. The LSTM model predicted stock prices better than all other models because it can handle both time-based and non-linear stock price movements. The model underwent further training after its maximum performance was confirmed because its design needed to process all available data from March 18 2025 until April 1 2025. The model predicted the closing stock price for 15 days which started from March 18 2025 until April 1 2025. The system achieved better results because it tracked the current market patterns together with the usual seasonal changes that drive stock market activity. The results were compared with actual results and it was established that the results were consistent with actual results in terms of reliability of this LSTM model in terms of actual and outcome curves.

3.1 Quantitative Performance Metrics

The models executed their training process using 80 percent of the data while their performance evaluation occurred through testing the remaining 20 percent of data which was divided into the test set. Table 2 gives the performance metrics of all models on the test data.

Table 2: Quantitative performance metrics (MAE, MSE, RMSE, and MAPE) for all five forecasting models evaluated on the 20% test set

Models	MAE	MSE	RMSE	MAPE
SARIMA	2.8803	12.3042	3.5077	0.2564
LSTM	1.4621	3.0479	1.7458	0.1117
Prophet	5.0385	28.0640	5.2975	0.3837
XGboost	1.9205	5.7707	2.4022	0.1471
Transformer	5.7222	36.6194	6.0514	0.4327

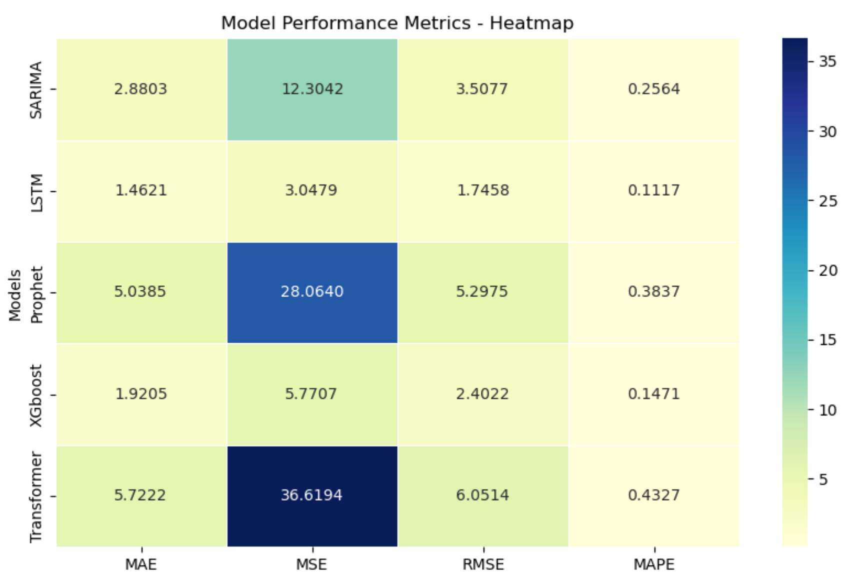

Fig. 8: Heatmap of performance metrics (MAE, MSE, RMSE, MAPE) across all five forecasting models

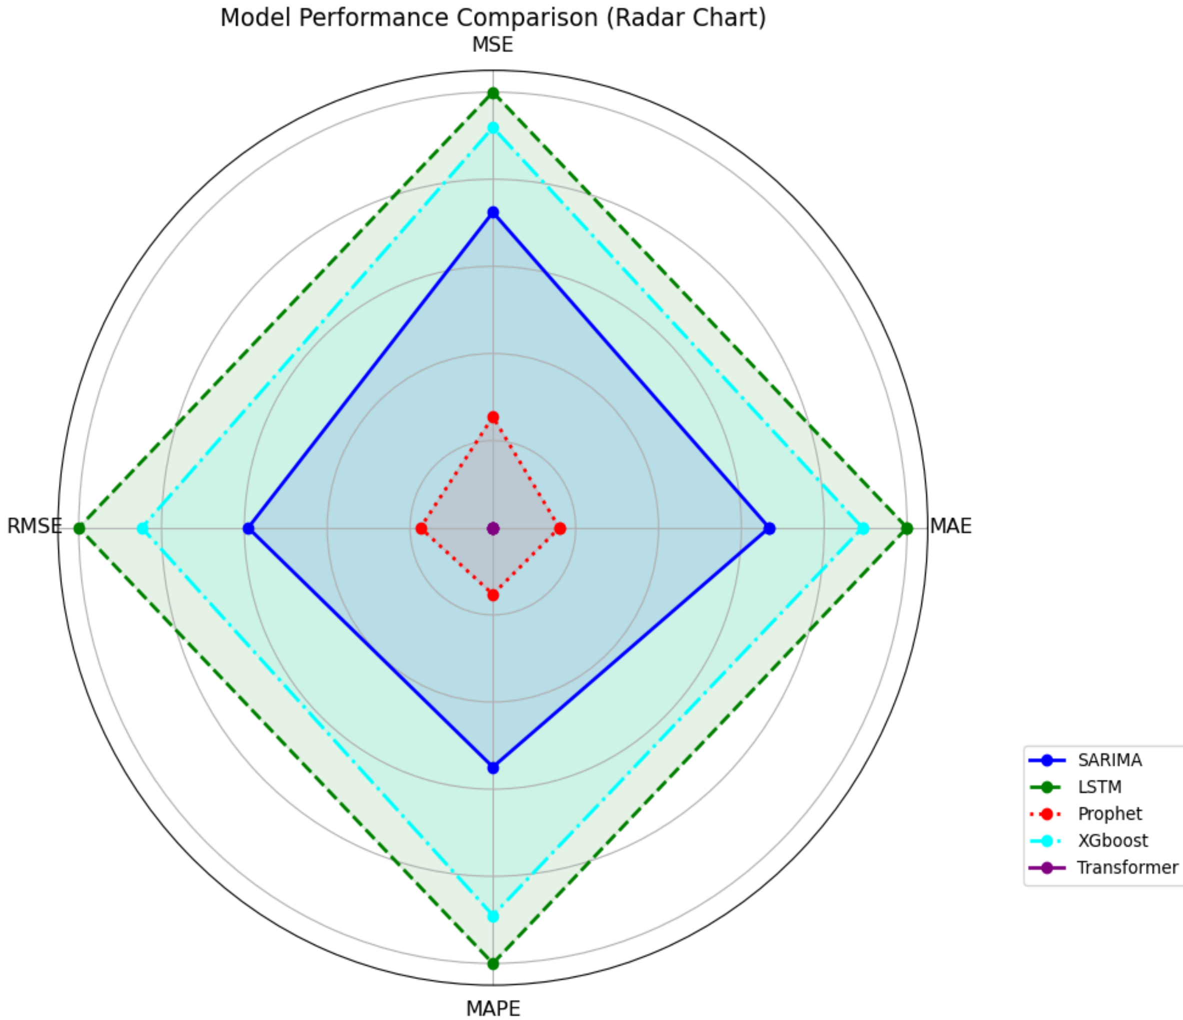

Fig. 9: Radar chart comparing the normalised performance metrics (MAE, MSE, RMSE, MAPE) of the five forecasting models

After comparing the parameters of performance criteria as mentioned in Table 2 above, the LSTM model seems to be more acceptable as it has shown the lowest value in the performance criteria of MAE: 1.4621, MSE: 3.0479, RMSE: 1.7458, and MAPE: 0.1117 or 11.17 percent. This is a clear illustration of high accuracy and reliability of the method to model accurately the seasonal and non-linear trends of the data. The remaining methods found more error (MAE: 2.8803, RMSE: 3.5077, MAPE:0.2564 or 25.64%), hence their challenges encountered in the process of addressing the non-linear patterns despite the assistance of the seasonal patterns. In the meantime, the XGBoost strategy had an intermediate performance (MAE: 1.9205, RMSE: 2.4022, MAPE:0.1471 or 14.71%), but not as accurate as the LSTM to allow the analysis because of the restrictions of the fixed variables. It is also valuable to mention based on the table above that the other method downloaded the worst performance because the criteria reported the highest levels of error (MAE: 5.7222, RMSE: 6.0514, MAPE:0.4327 or 43.27). This is attributable to the small sample size of the approach constituting of 733 samples; thus the poor performance of approach based on the attention mechanism presented by the approach because of the reliance on the small sample size represented by the approach in Table 2 and Figure 8 & Figure 9.

3.2 Key Findings

The findings indicate that LSTM is the best algorithm in the prediction of the close stock prices. Its ability to support the long dependency in the data as well as very low error rate and high explanatory values among others stood out. Although light, SARIMA and Prophet were the least efficient and, therefore, it is important to focus on the use of advanced techniques in the prediction of financial time-series. XGBoost and Transformers were good substitutes too but could not give pinpoint prediction as LSTM according to the nature of this dataset. Thus, it has been used in the prediction of the next 15 days closing prices by applying the LSTM model.

Table 3: LSTM-predicted daily closing prices for RTPower Ltd. over the 15-day forecast horizon

Date	Predicted Closing Prices
18-03-2025	8.532076
19-03-2025	8.482436
20-03-2025	8.401825
21-03-2025	8.263274
22-03-2025	8.100186
23-03-2025	7.873961
24-03-2025	7.671197
25-03-2025	7.482102
26-03-2025	7.285737
27-03-2025	7.091872
28-03-2025	6.901673
29-03-2025	6.733503
30-03-2025	6.560396
31-03-2025	6.385987
01-04-2025	6.214466

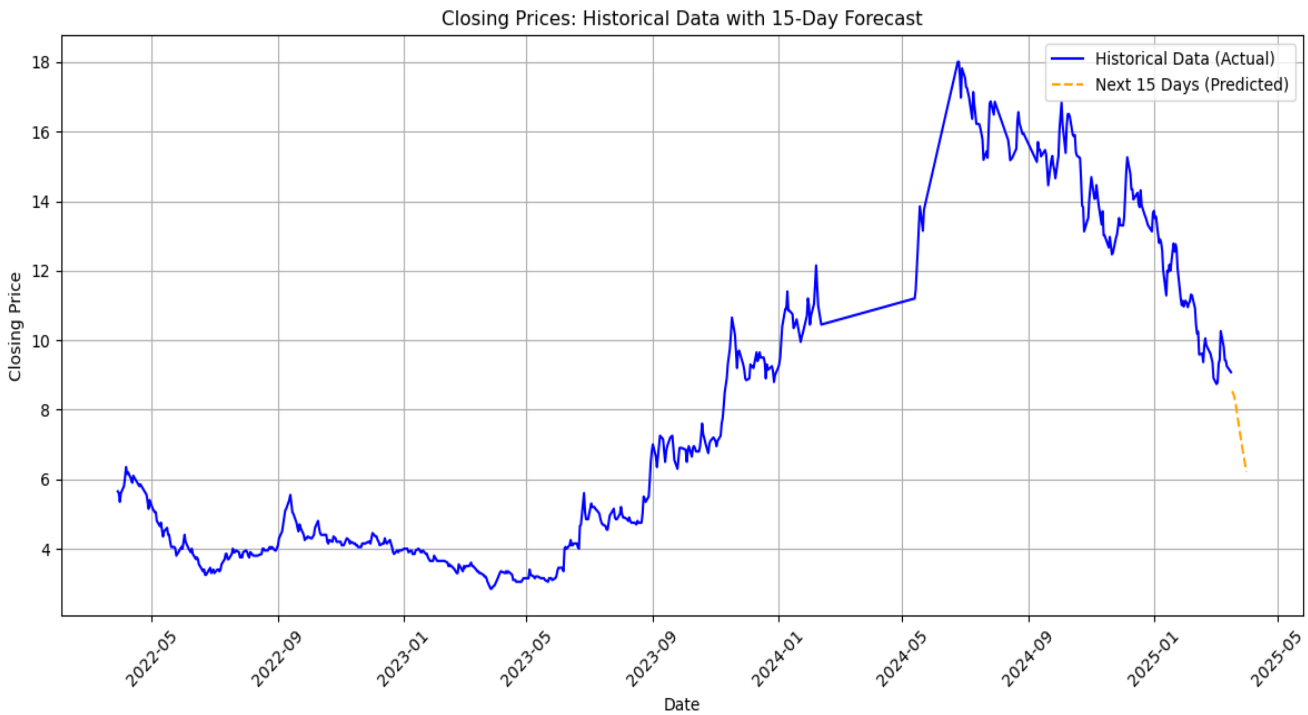

Fig. 10: Line plot showing the historical closing prices of RTPower Ltd. from March 29, 2022 to March 17, 2025, alongside the 15-day LSTM-predicted closing prices

Table 3 and figure 10 present the main information about the expected closing price of a stock in RTPower Ltd. between March 18, 2025 and April 1, 2025, and the previous performance between March 29, 2022, and March 17, 2025. The trend in the table of projected closing prices clearly indicates that there is a trend of falling, where the opening value of 8.532076 on March 18 and a closing value of 6.214464 on April 1 portray a decline of nearly 27 percent. The evidence also shows very drastic changes between, with values reaching a high of 8.263274 on March 21; this was quite a significant drop towards the end. This adverse forecast is observable following the previous trends depicted in this graph with the peak at nearly 18 in the middle of 2023 and the declining pace over the course of three years, almost until the beginning of 2025, showing nearly 9. It can be seen that this has been very volatile based on this data which represented nearly three years of previous trend with nearly sporadic increases in 2022-2023 that accumulate in nearly three years in a corrective cycle. This apparent decline however is supposed to be sustained as indicated in this research, though this decline to nearly 6.5 as of April 1st may indicate overestimation. This could fit the trend observed in the end of 2024 and early 2025, which would further support the validity of this research, but there is no notable upswing, in contrast to the previous ones, which appears to be conservative.

4. Discussion

The results highlight the robustness and supremacy of the LSTM model in predicting the next 15 days' closing stock prices, based on past experiences from March 29, 2022, to March 17, 2025. This section covers the implications of LSTM performance, relative advantages and disadvantages of individual models used, limitations of research, and potential avenues for practical application.

4.1 Interpretation of LSTM's Superiority

The high performance of the LSTM model is estimated by the MAE value of 1.4621, the RMSE value of 1.7458 and the MAPE value of 0.1117 and hence repeats with the fact that the LSTM model is able to consider the complex temporal dynamics and the consequent decline in closing stock values. This downwards trend may be observed to have a relative high of approximately 18 in mid-2023 up to approximately 9 on the date in March 2025. The structural flexibility of the LSTM model to account the long-term dependency patterns enables it to accommodate both gradual and abrupt changes in the values of stocks with a forecast to be reduced by 8.532076 on March 18, 2025, to 6.214464 on April 1, 2025. The consistency of the values on the test data set once again also emphasizes the ability of the LSTM algorithm to process the complexities of the data set of the stock market time-series data set, hence restating the possibility of deep learning algorithm to work well even when other algorithms may just happen to work out to be naive in predicting the stock closing value within the short run. The LSTM system enables stock traders and financial analysts to create precise predictions that extend up to 15 days into the future. The forecast period value decrease of 8.53-6.21 will help to establish which trading methods will enable to sell through the activation of signals that permit sales between 7.0-7.5 and through enhanced risk management that derives from their ability to accurately estimate future value drops. The model will be used in the late 2024 to March 2025 period because it matches the historical depreciation patterns which the market experienced during that time.

4.2 Limitations of the Study

Beyond the limitations specific to this study, it is worth acknowledging some of the broader challenges that come with using deep learning for stock price forecasting in general. Models that draw on a single source of data historical trading figures alone will always struggle to capture the full picture of how markets behave, especially when shifts in sentiment or broader economic mood are left out of the equation. There is also the question of how well these models travel across different industries and markets; many existing approaches have only been validated in narrow contexts, which raises genuine questions about how far their findings can be generalized.

This study is not exempt from either of these concerns. The model relies entirely on historical price data from one NSE-listed stock, and it does not account for external forces that can move markets in ways no price chart can anticipate — things like macroeconomic indicators, breaking news, or geopolitical developments. While the model holds up well when tested against historical data, whether it would perform with the same consistency in a live trading environment is a question that remains open. Future work would benefit greatly from weaving in richer data sources, such as real-time news feeds, social media sentiment, and broader economic signals.

It is also worth noting that deep learning models can generate misleading signals when the temporal context of a prediction is not carefully considered. For this reason, the 15-day forecast produced in this study should be treated as one input among many rather than a definitive guide to trading decisions, and practitioners are encouraged to interpret it with a healthy degree of caution [26-29].

4.3 Practical Applications and Future Directions

The fact that LSTM came out on top in this study is not just an academic result, it has real, tangible implications for how investors and analysts might approach short-term market decisions. Research has shown that portfolios guided by LSTM-based forecasts tend to deliver better daily returns, stronger cumulative gains, and lower volatility compared to portfolios put together without such guidance, which ultimately translates into more favorable risk-return outcomes in actual trading conditions.

The 15-day forecast generated in this study tells a fairly clear short-term story: a projected decline from 8.532076 to 6.214466. For traders paying attention, this kind of trajectory can serve as an early signal to consider selling in the 7.0–7.5 price range, allowing them to get ahead of what the model suggests could be a bearish stretch rather than reacting to it after the fact.

Looking further ahead, there is also a compelling case for extending the current model by pairing LSTM's price-tracking strengths with sentiment signals drawn from financial news. Studies combining both modalities have found improvements in generalizability, reduced model bias, and sharper forecasting efficiency compared to approaches that rely on price data alone suggesting this is a natural and worthwhile direction for future development.

That said, working with LSTM models in practice is not without its challenges. Getting the parameters right is a genuinely complex process, and poorly tuned models can underperform significantly, which is why the systematic hyperparameter tuning carried out in this study matters beyond just this particular dataset. More broadly, the comparative framework built here offers something practically useful in its own right — a structured basis for analysts to decide which model best fits the characteristics of their data and the time horizon they are working with.

5. Conclusion

The researcher conducted this study to predict stock prices which would occur during each day of a 15-day period which began on March 18 2025 and ended on April 1 2025. The dataset used for this study examined the period between March 29 2022 and March 17 2025 to evaluate LSTM model performance against SARIMA XGBoost Prophet and Transformer model performance. The LSTM model represents the most effective solution for this problem because it achieved an MAE measurement of 1.4621 and an RMSE measurement of 1.7458 and a MAPE measurement of 0.1117. The model demonstrates its capacity to analyze complete data sets, which enables it to track the downward trend that initiated after the peak in 2023 until the forecasted decline reaches 9 in March 2025. The current research demonstrates how advanced models can be utilized for predicting financial data during short-term periods. The system demonstrates its expected performance when measuring errors which reach their minimum value, while correctly predicting the price drop from 8.532076 on March 18, 2025, to 6.214464 on April 1, 2025. The model demonstrates its effective capacity to create accurate stock market predictions through its ability to process non-linear and sequential data. The LSTM model's accurate prediction of 15-day market trends throughout this study provides valuable insights which enable traders and investors to make selling decisions between 7.0-7.5 while foreseeing an upcoming bearish market trend. The model remains unresponsive to stock price fluctuations.

Acknowledgement

The author S. Aljohani would like to thank Prince Sultan University for paying the APC and for the support through the TAS research lab.

References

[1] Bodie, Z., Kane, A., & Marcus, A. J. (2014). *Investments* (10th ed.). McGraw-Hill Education.

[2] Kim, K. J. (2003). Financial time series forecasting using support vector machines. *Neurocomputing*, 55(1–2), 307–319.

[3] Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.

[4] Fama, E. F. (1991). Efficient capital markets: II. *The Journal of Finance*, 46(5), 1575–1617.

[5] Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.

[6] Brockwell, P. J., & Davis, R. A. (2016). *Introduction to time series and forecasting* (3rd ed.). Springer.

[7] Zhang, G. P. (2003). Time series forecasting using a hybrid ARIMA and neural network model. *Neurocomputing*, 50, 159–175

[8] Box, G., & Jenkins, G. M. (1976). Analysis: Forecasting and control. *San francisco*, *10*.

[9] Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785-794).

[10] Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015). Predicting stock market index using fusion of machine learning techniques. *Expert systems with applications*, *42*(4), 2162-2172.

[11] Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, *270*(2), 654-669.

[12] Li, Z., Yu, H., Xu, J., Liu, J., & Mo, Y. (2023). Stock market analysis and prediction using LSTM: A case study on technology stocks. *Innovations in Applied Engineering and Technology*, 1-6.

[13] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

[14] Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Philosophical transactions of the royal society a: mathematical, physical and engineering sciences*, *379*(2194).

[15] Khadka, S., Thapa, P., Sharma, P., Silwal, S., & KC, S. (2025). Do ensemble models always win? Comparative analysis of ML methods for stock price prediction. *Comparative Analysis of ML Methods for Stock Price Prediction (October 13, 2025)*.

[16] Shrouf, H., Shwedeh, F., Alzoubi, H. M., Aburub, F., & Joghee, S. (2026). Integrating Machine Learning Algorithms and Business Intelligence: Enhancing Decision-making in Food and Beverage Manufacturing Sector.

[17] Tan, C. L., Alzoubi, H. M., Yanamandra, R., Barakat, S., & Shwedeh, F. (2025). Supply Chain Disruption Prediction Using Machine Learning. *2025 International Conference on Business Intelligence for Technology Innovation (ICBITI)*, 1–8.

[18] Xiao, J., Deng, T., & Bi, S. (2024, August). Comparative analysis of LSTM, GRU, and transformer models for stock price prediction. In *proceedings of the international conference on digital economy,*

[19] Sarkar, A., & Vadivu, G. (2025). *An Advanced Ensemble Deep Learning Framework for Stock Price Prediction Using VAE, Transformer, and LSTM Model (2025)*.

[20] Hisam, M. W., Khan, K., Khan, M. S., Arora, V., & Ahmad, A. (2026). GDP Forecasting for Policy: Evaluating ARIMA, Exponential Smoothing, and XGBoost Models for Fiscal, Monetary, and Welfare Planning.

[21] Divisekara, R. W., Jayasinghe, G. J. M. S. R., & Kumari, K. W. S. N. (2020). Forecasting the red lentils commodity market price using SARIMA models. *SN Business & Economics*, *1*(1), 20.

[22] Zha, W., Liu, Y., Wan, Y., Luo, R., Li, D., Yang, S., & Xu, Y. (2022). Forecasting monthly gas field production based on the CNN-LSTM model. *Energy*, *260*, 124889.

[23] Samal, K. K. R., Babu, K. S., Das, S. K., & Acharaya, A. (2019, August). Time series-based air pollution forecasting using SARIMA and prophet model. In *proceedings of the 2019 international conference on information technology and computer communications* (pp. 80-85).

[24] Dar, A. A., Sanyal, S., Khan, K., Khan, M. S., Smerat, A., & Arora, V. (2025). Forecasting Retail Sales in the Indian Market: Integrating Time-Series and Machine Learning Approaches. *Ingenierie des Systemes d'Information*, *30*(11), 3029.

[25] Kwarteng, S., & Andreevich, P. (2024). Comparative analysis of ARIMA, SARIMA and Prophet model in forecasting. *Research & Development*, *5*(4), 110-120.

[26] Deswal, V., & Kumar, D. (2024). Hyperparameter tuning of the lstm model for stock price prediction. *International Journal of Intelligent Systems and Applications in Engineering*, *12*(4), 705-712.

[27] Vuong, P. H., Phu, L. H., Van Nguyen, T. H., Duy, L. N., Bao, P. T., & Trinh, T. D. (2024). A bibliometric literature review of stock price forecasting: From statistical model to deep learning approach. *Science progress*, *107*(1), 00368504241236557.

[28] Zhen, K., Xie, D., & Hu, X. (2025). A multi-feature selection fused with investor sentiment for stock price prediction. *Expert Systems with Applications*, *278*, 127381.

[29] Jagadesh, B. N., RajaSekhar Reddy, N. V., Udayaraju, P., Damera, V. K., Vatambeti, R., Jagadeesh, M. S., & Koteswararao, C. (2024). Enhanced stock market forecasting using dandelion optimization-driven 3D-CNN-GRU classification. *Scientific Reports*, *14*(1), 20908.

[30] Muhammad, D., Ahmed, I., Naveed, K., & Bendechache, M. (2024). An explainable deep learning approach for stock market trend prediction. *Heliyon*, *10*(21).