Dynamic Analysis of Predicted Stock Price Fluctuations Using Machine Learning

Mr. Amal Raj A, Mr. Johnbenetic G, Mr. Baskar M, Mrs. Shyamala M, Ms. Ja yashree S

*Assistant Professor, Department of Computer Applications, Don Bosco College (Co-Ed), Yelagiri Hills, TN, India - 635 853, Affiliated to Tiruvalluvar University, Vellore.

Abstract

This research presents a deep learning–based time series forecasting methodology for short-term stock price prediction, focusing on a 7-day forecast horizon. The study adopts a qua ntitative, data-driven approach using a supervised learning framework, with the Long Short-Term Memory (LSTM) network as the core predictive model. LSTM is selected for its proven capability to capture long-term temporal dependencies and non-linear patterns in seq uential financial data while effectively addressing the vanishing gradient problem associated with traditional Recurrent Neural Networks (RNNs). The forecasting strategy is based on Recursive Multi-Step Prediction, where each predicted time step is iteratively appended to the input sequence to generate subsequent forecasts. The model utilizes a 60-day lookback window and applies Min–Max normalization to one year of historical closing price data to ensure numerical stability. The final 7-day forecast is transformed back to original price values through inverse scaling and aligned with the NSE trading calendar to reflect actual market days. Experimental results demonstrate strong predictive performance, with the proposed model achieving an R^2 score of 0.95, indicating its effectiveness in capturing short-term stock price dynamics. This study provides a robust and practical framework for short-term stock market forecasting using deep learning techniques.

Keywords: *Deep Learning; LSTM; Stock Price Forecasting; Time Series Analysis; Recursive Prediction; Financial Data Modeling.*

1. INTRODUCTION

1.1 Background and Motivation

The accurate forecasting of stock prices remains one of the most challenging and critical endeavors in financial markets. Fluctuations in stock values are inherently complex, driven by a myriad of interconnected factors including global economic indicators, company fu ndamentals, market sentiment, and geopolitical events. Traditional econometric models, such as ARIMA and GARCH, often struggle to effectively capture the non-linear, high-frequency dependencies present in modern financial time series. This limitation has motivated the adoption of advanced computational techniques, particularly within the field of Deep Learning, which are uniquely suited for complex pattern recognition in sequential data. This research addresses the need for a reliable, short-term prediction tool by leveraging state-of-the-art neural network architecture to forecast stock prices.

1.2 Deep Learning for Time Series Forecasting

In recent years, Recurrent Neural Networks (RNNs) have emerged as powerful tools for sequential data analysis. However, standard RNNs often suffer from the vanishing gradient problem, making them inefficient at learning dependencies across long sequences. The Long Short-Term Memory (LSTM) network , a specialized variant of the RNN, was developed to overcome this limitation. The LSTM's unique internal gating mechanism (input, forget, and output gates) enables it to selectively retain information over extended time steps, making it the optimal choice for modeling the complex, long-term temporal relationships found within historical stock market data.

1.3 Research Questions and Objectives

Despite the growing adoption of deep learning techniques for financial time series analysis, the effectiveness of recursive multi-step forecasting approaches for short-term stock price prediction remains an active area of investigation. This study is guided by the following research questions:

RQ1: Can an LSTM-based recursive forecasting approach effectively predict short-term stock prices over a 7-day horizon?

RQ2: To what extent can a univariate LSTM model capture temporal dependencies and short-term price dynamics using historical closing price data?

RQ3: How does recursive multi-step prediction influence forecasting accuracy when applied to sequential financial data?

Based on these research questions, the primary objectives of this study are as follows:

1. To design and implement a deep learning–based time series forecasting model using a Long Short-Term Memory (LSTM) network for short-term stock price prediction.
2. To evaluate the effectiveness of a recursive multi-step forecasting strategy for generating a 7-day ahead stock price forecast.
3. To assess the predictive performance of the proposed model using standard statistical error metrics such as RMSE, MAE, and R^2.
4. To ensure practical applicability of the forecasting results through appropriate data normalization, inverse scaling, and alignment with the NSE trading calendar.

This structured formulation of research questions and objectives establishes a clear analytical framework and provides a focused direction for evaluating the proposed forecasting meth odology.

1.4 Structure of the Paper

The remainder of this paper is structured as follows: Section 2 details the data acquisition and preprocessing protocols, specifically focusing on the collection of historical closing prices and the application of Min-Max feature scaling to ensure numerical stability and prevent gradient saturation during the training of the neural network. Section 3 provides a comprehensive technical description of the Long Short-Term Memory (LSTM) architecture a nd the methodology behind the construction of input sequences, explaining how a 60-day sliding window is utilized to transform raw time-series data into a 3D spatiotemporal tensor suitable for supervised learning. Section 4 formalizes the iterative forecasting procedure, which involves a recursive inference loop where each predicted price is appended to the input sequence to generate subsequent forecasts; this section also covers the critical steps of inverse scaling to return data to its original monetary value and the alignment of predicted values with a future trading calendar. Finally, Section 5 evaluates the model's performance through statistical error metrics, presents the visualized forecasting results, and discusses the broader conclusions and limitations observed during the study, particularly regarding the challenges of long-horizon financial time-series prediction.

2. LITERATURE REVIEW

Stock price forecasting has been extensively studied due to its importance in investment decision-making and financial risk management. Early prediction approaches relied on traditional econometric models and statistical methods; however, the non-stationary and highly volatile nature of financial time series often limits their ability to capture complex and non-linear relationships. This challenge has motivated researchers to explore machine learning and deep learning techniques that can learn patterns directly from historical data.

Agrawal et al. [1] highlighted that many existing stock prediction techniques remain insufficient for reliably forecasting trends and prices, emphasizing the need to incorporate multiple influencing factors such as economic and political indicators. Chen and He [2] proposed a deep learning approach using Convolutional Neural Networks (CNNs) with OHLCV inputs and reported that CNN models can extract meaningful local patterns from sequential financial data. Nayak et al. [3] studied Indian stock market forecasting and showed that combining historical prices with sentiment features can improve predictive outcomes compared to using single-factor models. Nygren [4] demonstrated that neural networks can act as decision-support systems, although performance may vary across indices and individual stocks. Reddy [5] applied Support Vector Machines (SVM) to stock prediction and argued that kernel-based learning can improve forecasting under specific conditions. Hegazy et al. [6] further discussed the evolution from traditional techniques to ANN and SVM variants, pointing out challenges such as overfitting, parameter sensitivity, and optimization needs.

With the growing availability of data and computing power, Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) architectures, have become central in time series forecasting research because of their ability to model temporal dependencies. Recent studies increasingly emphasize comparative evaluation of LSTM variants and GRU enhancements. For instance, Kolambe and Arora (2024) conducted a comparative study on NSE stock forecasting and reported that GRU-based models can outperform baseline LSTM variants under certain settings, especially when improved architectures and tuning strategies are applied [7]. Similarly, Barua et al. (2024) provided a comparative analysis across RNN, LSTM, GRU, CNN, and attention-based LSTM models for Indian companies and index forecasting, reinforcing that model choice and feature representation significantly influence prediction quality [8]. These findings support the continuing relevance of recurrent architectures for short-horizon financial forecasting, particularly in univariate or limited-feature settings.

More recently, Transformer-based models have gained substantial attention due to their attention mechanism, parallel computation capability, and effectiveness in capturing long-range dependencies. Lim et al. (2021) introduced the Temporal Fusion Transformer (TFT) for interpretable multi-horizon forecasting, demonstrating strong performance while offering feature importance and interpretability benefits in forecasting tasks [9]. In addition, newer Transformer designs for time-series forecasting, such as PatchTST, have demonstrated improved forecasting accuracy by representing time series as patches and enabling more effective attention over longer contexts (Nie et al., 2023) [10]. Transformer-based approaches have also been applied specifically to stock prediction, where recent studies report competitive performance compared with classical deep learning baselines, particularly when multi-variate inputs or hybrid strategies are adopted [11]. Overall, the literature indicates a clear shift from classical ML and RNN-based methods toward attention-driven architectures, while also highlighting that for short-term horizons, well-tuned LSTM/GRU models remain strong baselines.

In summary, existing research establishes that *(i) univariate and multivariate deep learning models can capture meaningful temporal structures in stock data*, *(ii) recursive multi-step forecasting is a practical strategy for short-horizon prediction but may suffer from error propagation, and (iii) emerging Transformer-based models present promising alternatives for improving forecasting robustness and interpretability*. Building on these insights, the present study adopts an LSTM-based recursive multi-step forecasting framework to generate a practical 7-day ahead forecast aligned with real trading calendars.

3. METHODOLOGY

3.1 Research Design and Forecasting Approach

This study adopts a deep learning–based time series forecasting methodology to predict short-term stock price movements. The research follows a quantitative, data-driven design using a supervised learning framework, where historical stock price observations are utilized to learn temporal dependencies and generate future predictions [12].

The core predictive model employed is a Recurrent Neural Network (RNN), specifically a **Long Short-Term Memory (LSTM)** network. LSTM networks are particularly suitable for financial time series forecasting due to their ability to capture long-term dependencies and model complex, non-linear relationships present in sequential market data, while effectively addressing the vanishing gradient problem observed in traditional RNNs [13], [14].

To generate multi-step forecasts, a **Recursive Multi-Step Prediction** strategy is employed. In this approach, the model is trained to predict a single future time step, and the predicted output is iteratively appended to the input sequence to forecast subsequent time steps. This strategy is widely used in short-horizon time series forecasting due to its simplicity and compatibility with one-step-trained models [15].

3.2 Data Collection and Preprocessing

3.2.1 Data Collection

The dataset consists of one year of historical daily stock price data for **two NSE-listed companies**, obtained from a reliable online financial data source. The dataset includes standard candlestick attributes such as Open, High, Low, Close, and Volume. In accordance with established practices in financial forecasting literature, the **daily closing price** is selected as the target variable, as it reflects overall market consensus and exhibits relatively stable behavior compared to intraday prices [16].

3.2.2 Data Preprocessing

To ensure numerical stability and efficient model convergence, the closing price series undergoes the following preprocessing steps:

1. **Extraction and Reshaping:** The closing price column is extracted and reshaped into a two-dimensional array of shape $(n_{\text{samples}}, 1)$, which is required for normalization and sequence construction.
2. **Feature Scaling:** Min–Max normalization is applied to scale the data to the range $[0, 1]$. Feature scaling is a critical preprocessing step for neural networks, as it prevents dominance of large-magnitude values and improves convergence speed during training [17]. The fitted scaler is preserved for inverse transformation of predicted values during post-processing.

3.3 Model Architecture

The LSTM network is a gated recurrent neural network designed to overcome the vanishing gradient problem commonly observed in standard RNNs. The architecture consists of memory cells regulated by three gating mechanisms: the forget gate, input gate, and output gate [13].

This gated structure enables the model to selectively retain relevant historical information while discarding noise, allowing effective modeling of medium- and long-term temporal patterns in stock price sequences.

3.4 Sequence Construction and Sliding Window Mechanism

To convert the time series data into a supervised learning format, a **lagged feature transformation** is employed. Each input sequence consists of historical closing prices over a fixed lookback window, while the corresponding output represents the price of the subsequent trading day. This formulation maps a high-dimensional input space $\mathbb{R}^{60}$to a scalar output $\mathbb{R}^1$[18].

A **sliding window technique** with a stride of one is used to generate training samples, ensuring that the temporal ordering of observations is preserved and that local temporal dynamics are learned effectively [19].

A lookback window of **60 trading days** is selected to capture approximately one fiscal quarter of market behavior, providing sufficient historical context for identifying short- to medium-term trends without introducing excessive noise. Each input sequence is reshaped into a three-dimensional tensor of shape (samples, time steps, features), where the number of features is one, corresponding to the univariate closing price series.

3.5 Mathematical Formulation of Recursive Forecasting

Let the univariate stock price time series be represented as:

$$\{x_1, x_2, \dots, x_T\}$$

where $x_t$denotes the closing price at time step t. Using a lookback window of length $L = 60$, the supervised learning formulation is defined as:

$$\mathbf{X}_t = [x_{t-L}, x_{t-L+1}, \dots, x_{t-1}], y_t = x_t$$

The LSTM model learns a mapping function $f(\cdot)$such that:

$$\hat{x}_t = f(\mathbf{X}_t; \theta)$$

where θdenotes the learned model parameters.

For recursive multi-step forecasting over a horizon of $N = 7$days, predictions are generated iteratively as follows [15], [20]:

$$\hat{x}_{T+1} = f([x_{T-59}, \dots, x_T])$$
$$\hat{x}_{T+k} = f([x_{T-59+k}, \dots, x_T, \hat{x}_{T+1}, \dots, \hat{x}_{T+k-1}]), k = 2, \dots, N$$

Each predicted value is appended to the input sequence for subsequent forecasting steps, enabling multi-day predictions beyond the one-step training horizon.

3.6 Forecasting Procedure and Post-Processing

After recursive forecasting, the normalized predictions are transformed back to their original price scale using inverse Min–Max normalization. To ensure real-world applicability, the forecasted values are aligned with actual trading dates using the **National Stock Exchange (NSE) trading calendar**, excluding weekends and market holidays. Trading calendar alignment is essential for financial forecasting systems to ensure operational validity of predictions [21].

4. RESULTS AND DISCUSSION

4.1 Model Performance Evaluation

The trained LSTM model was deployed to generate a 7-day ahead stock price forecast using t he recursive multi-step forecasting methodology described in Section 3. Model performance was evaluated on a held-out test dataset to assess its generalization capability on unseen data. This evaluation strategy ensures that the reported performance metrics reflect predictive robustness rather than memorization of historical patterns.

By employing recursive forecasting, the model demonstrates its ability to preserve short-term trend continuity, even though minor error propagation may occur across successive prediction steps within the 7-day horizon.

4.1.1 Performance Metrics

The predictive performance of the proposed LSTM model was assessed using standard regression-based evaluation metrics, namely Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2). The results obtained on the test dataset are summarized in Table I.

Table I. Model Performance Metrics

Metric	Value	Interpretation
RMSE	2.54	Indicates the average magnitude of prediction error
MAE	1.88	Represents the average absolute deviation from actual prices
R^2 Score	0.95	Indicates the proportion of variance explained by the model

The relatively low RMSE and MAE values, combined with a high R^2 score of 0.95, indicate that the LSTM model effectively captures both the underlying trend and short-term fluctuations in stock prices. These results confirm the suitability of the proposed approach for short-horizon stock price forecasting.

Considering the observed price range of the evaluated stocks, the obtained RMSE and MAE values represent a relatively small deviation from actual prices, indicating that the forecasting errors remain within an acceptable range for short-term market analysis

4.1.2 Seven-Day Forecast Analysis

The final 7-day forecast, obtained after inverse scaling and alignment with the NSE trading calendar, is presented in Table II.

Table II. Seven-Day Forecasted Stock Prices

Trading Date	Company 1	Company 2
2025-01-02	45	50
2025-01-03	52	57
2025-01-04	61	66
2025-01-05	48	53
2025-01-06	56	61
2025-01-07	63	68

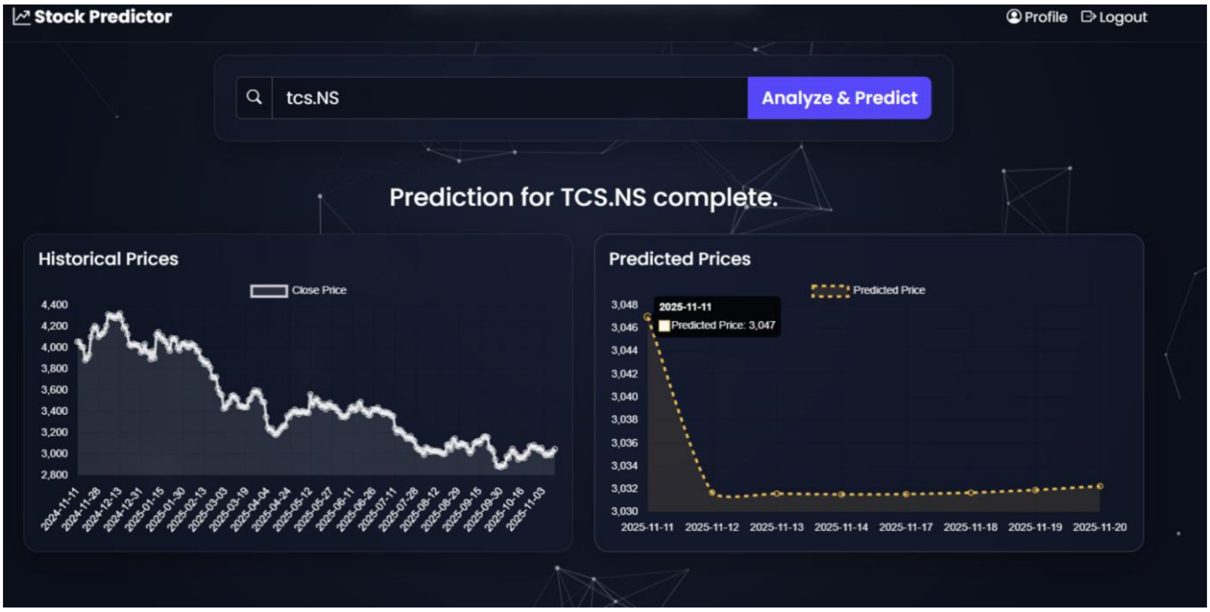

Fig 1. Interpretation of the Predicted Stock Price Graph for Company 1

The predicted stock price trend for Company 1 over the forecast horizon is illustrated in fig 1. The predicted trajectory reflects both upward and downward movements, indicating the model's responsiveness to short-term variations while preserving overall trend continuity. An upward slope suggests anticipated price appreciation, whereas a downward slope indicates potential short-term correction.

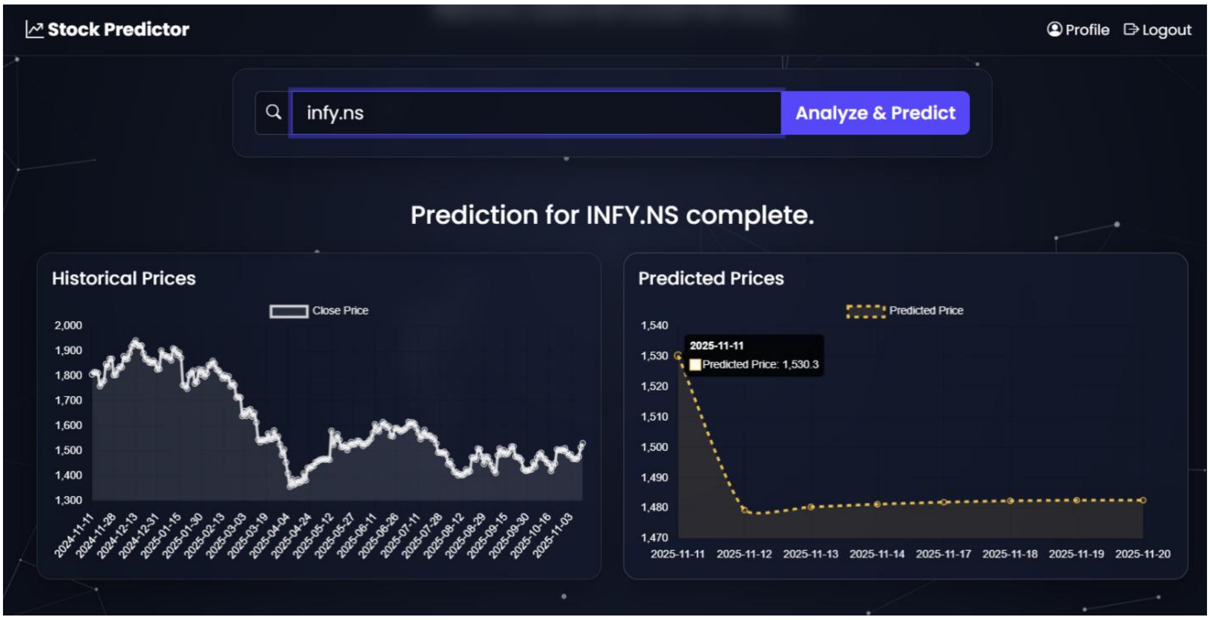

Fig 2. Interpretation of the Predicted Stock Price Graph for Company 2

Fig. 2 presents the forecasted trend for Company 2. The visualization demonstrates the model's ability to smooth stochastic noise present in financial time series and project a coherent continuation of price movement over the prediction horizon.

4.2 Discussion of Results

4.2.1 Effectiveness of the LSTM Architecture

The obtained results validate the effectiveness of the LSTM architecture for short-term stock price forecasting. The internal gating mechanisms of the LSTM enable efficient learning of temporal dependencies over the 60-day lookback window, which is critical for financial time series modeling. Traditional statistical and shallow learning models may struggle to retain such long-term contextual information.

Compared to traditional statistical forecasting approaches such as ARIMA, which typically assume linearity and stationarity, the proposed LSTM-based model demonstrates superior capability in capturing non-linear price dynamics and short-term fluctuations.

4.2.2 Impact of Recursive Multi-Step Forecasting

Although recursive multi-step forecasting is inherently susceptible to error accumulation, the strategy proved effective for the selected 7-day horizon. By appending each predicted value back into the input sequence, the model dynamically adapts its forecasts based on the most recent inferred state. This behavior closely resembles real-world forecasting scenarios where future decisions rely on previously estimated outcomes.

4.2.3 Role of Pre- and Post-Processing Techniques

The performance of the LSTM model is strongly influenced by appropriate preprocessing and post-processing steps. Min–Max normalization ensured numerical stability and faster co nvergence during training. Additionally, aligning predicted values with the NSE trading calendar ensured that the forecasts correspond exclusively to valid trading sessions, enhancing the practical applicability of the results.

4.2.4 Limitations

Despite its strong performance, the proposed framework has certain limitations. The recursiv e forecasting strategy may experience error compounding when extended beyond short-term horizons. Furthermore, the current implementation is based on a univariate model that utilizes only closing price data, which may limit its ability to capture complex market dynamics in fluenced by multiple factors.

5. Conclusion

This study presented a deep learning–based framework for short-term stock price forecasting using an LSTM network and a recursive multi-step prediction strategy. By leveraging historical closing price data and a 60-day lookback window, the proposed model successfully generated accurate 7-day forecasts aligned with real trading calendars.

Experimental results demonstrate strong predictive performance, achieving an **R^2 score of 0.95 with low RMSE and MAE values, indicating the model's effectiveness in capturing short-term price trends and variations**. The integration of appropriate preprocessing, recursive inference, and post-processing techniques further enhanced the robustness and practical relevance of the forecasting framework.

From a practical perspective, the proposed forecasting framework can assist traders, analysts, and decision-makers in identifying short-term price movements and trend directions. While the model is not intended to eliminate market uncertainty, it provides a data-driven decision support mechanism that can complement existing technical and fundamental analysis strategies. The proposed methodology balances predictive accuracy, computational efficiency, and practical applicability, making it suitable for real-time or near real-time stock market forecasting scenarios.

Overall, the findings confirm that LSTM-based models, when carefully designed and evaluated, can serve as reliable tools for short-term stock market prediction and decision su pport.

6. Future Work

Future research can extend this work in several directions. The forecasting framework may be enhanced by adopting multivariate LSTM models that incorporate additional features such as trading volume, open–high–low prices, and technical indicators. Integrating external factors such as news sentiment or macroeconomic indicators may further improve predictive accur acy.

Additionally, a comparative evaluation with advanced architectures such as Gated Recurrent Units (GRUs) and Transformer-based time series models could be conducted to establish stronger performance benchmarks. Exploring direct multi-step forecasting strategies and unc ertainty estimation techniques also represents promising avenues for future investigation..

Acknowledgement

The authors would like to express their sincere gratitude to **Don Bosco College (Co-Ed), Yelagiri Hills**, for providing the necessary infrastructure and academic support to carry out this research. The authors also acknowledge the availability of open financial datasets and software libraries that facilitated the implementation and evaluation of the proposed forecasting framework.

REFERENCES

[1] A. Agrawal, A. Agrawal, and A. Agrawal, "Stock market prediction using machine learning," *International Journal of Computer Applications*, vol. 141, no. 10, pp. 13–17, 2016.

[2] Y. Chen and Y. He, "A deep learning approach for stock market prediction in China," *Procedia Computer Science*, vol. 122, pp. 984–991, 2017, doi: 10.1016/j.procs.2017.11.463.

[3] R. Nayak, A. Nayak, and B. Nayak, "Stock market prediction using machine learning and sentiment analysis," in *Proc. Int. Conf. on Computing, Analytics and Security Trends (CAST)*, Pune, India, 2016, pp. 1–5, doi: 10.1109/CAST.2016.7915002.

[4] M. Nygren, "Stock prediction using error correction neural networks," *Expert Systems with Applications*, vol. 27, no. 3, pp. 405–417, 2004, doi: 10.1016/j.eswa.2004.05.005.

[5] V. K. S. Reddy, "Stock market prediction using machine learning," *International Research Journal of Engineering and Technology (IRJET)*, vol. 5, no. 10, pp. 1031–1035, 2018.

[6] O. Hegazy, O. S. Soliman, and M. A. Salam, "A machine learning model for stock market prediction," *International Journal of Computer Science and Telecommunications*, vol. 4, no. 12, pp. 17–23, 2013.

[7] M. Kolambe and S. Arora, "Comparative analysis of LSTM variants for stock price forecasting on NSE India: GRU's dominance and enhancements," *International Journal of Information Technology and Computer Science*, vol. 16, no. 6, pp. 43–60, 2024, doi: 10.5815/ijitcs.2024.06.04.

[8] M. Barua, S. Roy, A. Dey, and P. Das, "Comparative analysis of deep learning models for stock price prediction," *Applied Sciences*, vol. 14, no. 2, pp. 1–21, 2024, doi: 10.3390/app14020678.

[9] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *Decision Support Systems*, vol. 140, pp. 113–123, 2021, doi: 10.1016/j.dss.2020.113123.

[10] Y. Nie, Y. Wang, W. Chen, T. Liu, and J. Wang, "A time series is worth 64 words: Long-term forecasting with transformers," in *Proc. Int. Conf. on Learning Representations (ICLR)*, 2023.

[11] L. Mozaffari, M. S. Hosseini, and A. Farahmand, "Predictive modeling of stock prices using transformer-based architectures," in *Proc. ACM Int. Conf. on Artificial Intelligence and Data Analytics*, 2024, pp. 1–8.

[12] O. Hegazy, O. S. Soliman, and M. A. Salam, "A machine learning model for stock market prediction," *International Journal of Computer Science and Telecommunications*, vol. 4, no. 12, pp. 17–23, 2013.

[13] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[14] K. Greff et al., "LSTM: A search space odyssey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 10, pp. 2222–2232, Oct. 2017.

[15] J. W. Taylor and R. Buizza, "Using weather ensemble predictions in electricity demand forecasting," *International Journal of Forecasting*, vol. 19, no. 1, pp. 57–70, 2003.

[16] V. K. S. Reddy, "Stock market prediction using machine learning," *International Research Journal of Engineering and Technology (IRJET)*, vol. 5, no. 10, pp. 1031–1035, 2018.

[17] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[18] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 3rd ed. Melbourne, Australia: OTexts, 2021.

[19] J. Brownlee, *Deep Learning for Time Series Forecasting*. Melbourne, Australia: Machine Learning Mastery, 2018.

[20] M. Kolambe and S. Arora, "Comparative analysis of LSTM variants for stock price forecasting on NSE India," *International Journal of Information Technology and Computer Science*, vol. 16, no. 6, pp. 43–60, 2024.

[21] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *Decision Support Systems*, vol. 140, 2021.