LSTM-Based Temporal Analysis of Nifty 50: Accuracy Dynamics across Varied Time Frames

Priyanka Dash[1], Jyotirmaya Mishra[1], Suresh Dara[2*],

[1] Department of Computer Science and Engineering, GIET University, Odisha-765022, India
[2]School of Computer Science and Engineering, VIT-AP University, Andhra Pradesh-522237, India.
[1]priyankadash2018@gmail.com, [2]jyoti@giet.edu,
*[3]darasuresh@live.in

Abstract. Given the inherent noise and volatility in financial markets, accurate stock market price prediction is still a difficult task. In this work, the Investigation of Long Short-Term Memory (LSTM) models for stock price prediction in the face of varied levels of noise in various historical data time frames is carried out. Because of LSTM has reputation for capturing long-term dependencies, and these (LSTMs) are used to navigate the complex and noisy world of stock market dynamics. Several financial indicators and past stock prices are included in this set, which is pre-processed to account for the noise that is always present in financial data. The data is arranged temporally over a variety of intervals, from daily to monthly, which enables one to examine the flexibility of the LSTM model over a range of time periods. The study measures the predictive accuracy, precision, and recall of the LSTM model by methodically analysing its performance in the presence of data noise. The robustness of the model in both short- and long-term prediction situations receives particular consideration. Moreover, the influence of combining different technical indicators and outside variables is investigated to determine how well they work to reduce noise and improve forecast accuracy. This work recognises and addresses the ubiquitous noise in financial datasets, adding to the sophisticated knowledge of LSTM models in the context of stock market prediction. The results of this study demonstrate the LSTM are effective for stock price prediction the LSTM model has yielded an R2-score of 1 and Root Mean Squared Error (RMSE) value of 27.93.

Keywords: Stock Market Prediction, Long Short-Term Memory (LSTM), Data Noise, Financial markets, Long-term dependencies

1. Introduction

In the world of trading during the day, where investors aim to profit from brief price changes, real-time price forecasts are essential for developing successful trading plans. The difficulty stems from the stock market's intrinsic complexity, which is noisy, complicated, dynamic, volatile, non-parametric, and non-stationary. Stock forecasting has seen success with machine learning models, especially those built on the Long Short-

K. R. Madhavi et al. (eds.), *Proceedings of the International Conference on Computational Innovations and Emerging Trends (ICCIET 2024)*, Advances in Computer Science Research 112,
https://doi.org/10.2991/978-94-6463-471-6_111

Term Memory (LSTM) architecture. To account for the complexities of the always changing market conditions, it is necessary to adjust their hyper parameters using the most recent market data.Machine learning models are typically trained and evaluated in batches to speed up learning and facilitate error correction. For intraday stock predictions, however, a paradigm change is necessary. To guarantee improved accuracy in real-time predictions, models should learn continuously and forecast for each occurrence separately rather than the full batch. Using two unique learning approaches incremental learning and offline–online learning this research presents a fresh technique based on the LSTM model. The model is updated continually in incremental learning as it gets the next instance of the stock from the live market stream. On the other hand, offline-online learning entails retraining the model following every trading session to make sure it takes into account the most recent intricacies in market data.

2. Recent Studies

Stock price prediction plays a very important role for variety of reasons like Investment Decision making, Risk Mitigation, Portfolio Management, Financial Planning and Market Trend Analysis etc.. In this direction [1] examines the growing tendency of people buying stock market investments as a result of economic growth.
The research suggests a novel way for predicting the closing price of stocks for the following day, the CNN-BiLSTM-AM method, in light of the volatility of stock markets and the complexity of factors impacting stock prices. Convolutional neural networks (CNN), bi-directional long short-term memory (BiLSTM), and an attention mechanism (AM) are all integrated in this method.

The CNN is used to extract features from the input data, the BiLSTM uses these features to forecast the closing price of the stock the next day, and the attention mechanism records the influence of feature states at various points in the past with the goal of improving prediction accuracy. Using data from the Shanghai Composite Index over 1000 trading days, the study compares the suggested method with seven alternative strategies. [2] Classifies research based on several deep learning techniques, such as CNN, LSTM, DNN, RNN, Reinforcement Learning, and others like Hybrid Attention Networks and Wavenet, using articles from the Digital Bibliography & Library Project (DBLP) database. Each study's variables, models, and dataset characteristics are reviewed in the paper, and the findings are presented using important performance measures like accuracy, Sharpe ratio, MSE, RMSE, MAPE, and MAE. The survey notes the rising popularity of current models that fuse LSTM with other methods, such DNN, and emphasizes the performance and potential benefits of reinforcement learning along with other deep learning techniques. The conclusion highlights how deep learning-based financial modeling techniques have become exponentially more popular in recent years. [3] tackles the difficulty of forecasting the development trends in financial operations because financial data is ambiguous, complex, and incomplete. Financial data are time-dependent and nonlinear, with multiple associated factors influencing them. In contrast to conventional machine learning methods, Deep Neural Networks (DNNs) [4,5], which combine the advantages of deep learning and neural networks, are suggested as a way to approach nonlinear challenges in financial data processing more successfully. The financial product price data is treated in the paper as a one-dimensional series that is produced by projecting a chaotic system with several components into the temporal dimension [6]. The price series is reconstructed using the time series phase-space reconstruction (PSR) technique. Based on the PSR approach, a DNN-based prediction model [7] is created that incorporates long- and short-term memory networks (LSTMs) [8] for deep learning. The model predicts stock prices and is used to test its performance against other prediction models over different time periods and stock indexes. The comparison's findings show that the suggested prediction model—which combines DNNs and LSTMs—predicts more accurately than the other models, demonstrating its efficacy in predicting stock values in the context of intricate financial data.

3. Proposed methods

One important way to handle problems related to stock price prediction is to use the specialised recurrent neural network (RNN) architecture known as Long Short-Term Memory (LSTM). It is especially useful in financial forecasting applications because of its sturdy construction, which enables it to capture complex patterns and correlations in time series data. The cell state (Ct) and hidden state (ht) are the fundamental components of the LSTM architecture. During the LSTM chain, the cell state serves as the network's long-term memory. It is possible to add or delete information from the cell state selectively using a gating mechanism. The output for forecasts or for propagation to other time steps is the hidden state, which stands for short-term memory.

Important elements of LSTM operation in the context of stock price forecasting consists

Input Gate Assesses if the input data that will be incorporated into the cell state is relevant. The input gate generates values between 0 and 1, which represent the significance of the input data, using a sigmoid layer.

Forget Gate Selects the data from the cell state that should be ignored. A sigmoid layer is used, and values between 0 and 1 indicate whether information is retained (1) or omitted (0).

Cell State Update uses data from the input gate and forget gate to update the cell state. Tanh layers help to modify the cell state by producing a vector of new candidate values.

Output Gate uses the updated cell state to determine the next concealed state. The cell state, which has been altered by the tanh layer, is multiplied by the output of the sigmoid layer, which determines which portions of the cell state should be output.

At practice, LSTMs have shown to be successful at forecasting changes in stock prices, which helps investors, financial analysts, and algorithmic traders make well-informed choices. The sequential structure of stock price data allows LSTMs to recognize trends, seasonality, and other pertinent characteristics that help produce forecasts that are accurate. In conclusion, the LSTM architecture is a particularly effective instrument for stock price prediction. Because of their capacity to handle vanishing gradient problems and efficiently simulate temporal dependencies, long short-term memory banks (LSTMs) are a vital component in the creation of complex financial market forecasting models

4. Experimental Results and Discussions

The depicted line graph in Fig. 1 illustrate the fluctuating dynamics of a stock's low and open values over the period from February 2015 to February 2024. The x-axis chronicles the progression of time, while the y-axis denotes the stock price, ranging from approximately 8,000 in February 2015 to around 22,000 in February 2024. The general trend of the stock is one of growth, signifying an overall increase in value over the specified timeframe. However, the visual narrative also unveils sporadic instances of volatility, characterized by sharp declines in the stock price followed by subsequent recoveries. Notably, the stock's volatility appears to have intensified in more recent years compared to the earlier part of the depicted timeline. This observation suggests that the market has experienced heightened turbulence, potentially influenced by various economic or market-specific factors impacting the stock's performance. Overall, the graphs provide a comprehensive overview of the stock's price movements, capturing both its upward trajectory and the periodic challenges posed by increased volatility in the later years.

The presented line graph in Fig 2, arraying the close and high values of a stock from February 2015 to February 2024, offers valuable insights into the stock's performance over time. Notably, there is a discernible upward trend in the stock price, exemplified by the closing price reaching approximately ₹22,000 in February 2024, a substantial increase from the starting point of around ₹8,000 in February 2015. However, the graph also highlights periods of volatility, notably in December 2018 and March 2020, where the stock experienced sharp declines. Encouragingly, the stock price consistently rebounded from these downturns, showcasing resilience in its overall trajectory. A noteworthy observation is the heightened volatility in recent years, particularly evident in the more erratic lines during 2023 and 2024 compared to the relatively smoother trends in earlier years. This increased choppiness indicates a greater degree of market turbulence, suggesting that the stock has become more susceptible to fluctuations in the later years of the depicted timeline.

The graph further reveals that the highest closing price occurred in February 2024, with the stock concluding at around ₹22,000. Conversely, the lowest closing price is recorded in March 2020, when the stock closed at approximately ₹8,000. These extreme points underscore the range of market movements and emphasize the stock's capacity for both significant gains and recoveries from notable setbacks throughout the depicted period.

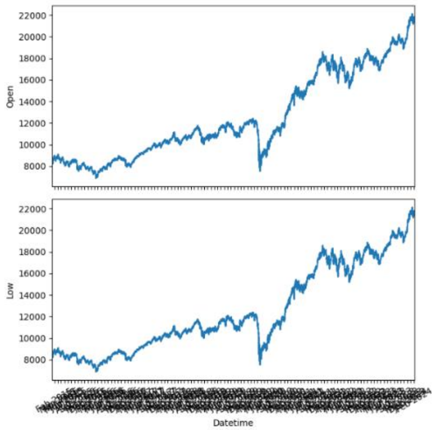

Fig. 1 Trend Analysis of Low and Open against Data

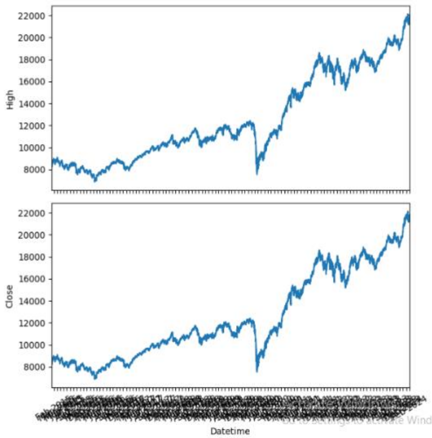

Fig. 2 Trend of Close and High prices of Stock

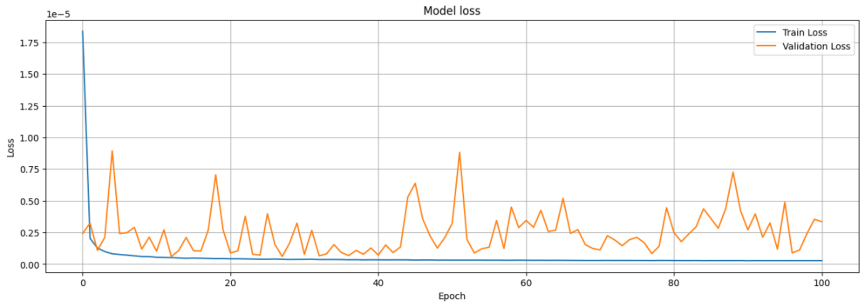

Fig. 3. Training and Validation Loss Curve

In Fig 3, the provided values represent the loss and validation loss obtained during the training of a machine learning model. The provided observations pertain to training and validation losses in a computational model. The training loss is remarkably low, measuring at 2.78E-07, indicating that the model performed exceptionally well in minimizing errors during the training phase. This suggests a high level of accuracy and efficiency in learning the patterns within the training dataset.

On the other hand, the validation loss, though slightly higher at 3.35E-06, remains relatively low. Validation loss is crucial for assessing how well the model generalizes to unseen data. The close proximity of the validation loss to the training loss indicates that the model has not suffered from overfitting, maintaining its effectiveness in making accurate predictions on new, unseen data.

The provided observations present various metrics evaluating the performance of a predictive model. The Mean Absolute Error (MAE) is calculated at 16.64, providing an average measure of the absolute differences between predicted and actual values. A lower MAE indicates better accuracy, and the reported value suggests a relatively moderate level of prediction error. The Mean Squared Error (MSE) is 779.92, reflecting the average of squared differences between predicted and actual values. This metric tends to magnify larger errors, resulting in a higher value compared to MAE. The Root Mean Squared Error (RMSE), derived from the square root of MSE, is 27.93, serving as a more interpretable measure. The RMSE indicates the average magnitude of prediction errors and, in this case, signifies a moderate level of deviation from actual values. The R-squared (R2) value of 1 is a perfect score, suggesting that the model explains 100% of the variance in the dependent variable. While a high R2 is desirable, it's essential to ensure it's not a result of overfitting. Further scrutiny of the model's generalization to new data is advised. The Mean Absolute Percentage Error (MAPE) is an impressively low 0.08%, indicating a minimal average percentage difference between predicted and actual values. The Median Absolute Percentage Error (MDAPE) is even lower at 0.05%, emphasizing the robustness of the model's predictions across a variety of scenarios.

Training and Test Sets: The data is split into training and test sets for each time interval. The training set typically contains a larger proportion of the data compared to the test set. This division allows the model to learn patterns from the training data and evaluate its performance on the test data.

For instance, in minute1, the training set comprises 647,829 rows, while the test set has 161,969 rows. This pattern holds for other intervals as well, with a similar split observed in decreasing magnitudes.

The division into training and test sets follows a logical progression across different time intervals, ensuring a representative distribution for model training and evaluation.

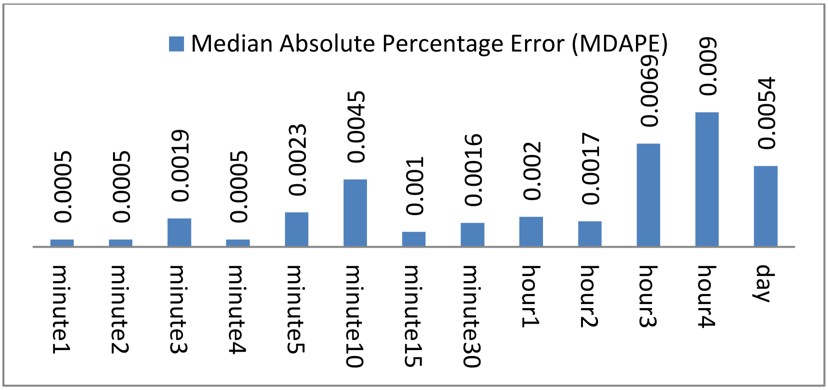

Fig.4. MDAPE for Different Data input to model

Fig. 4. shows the median absolute percentage error for different data inputs to model. In this dataset, the MDAPE values for shorter time intervals, such as minutes and hours, generally appear quite low, indicating accurate predictions with small percentage errors. However, as the forecasting horizon extends to a day, the MDAPE values increase, suggesting a larger central tendency of percentage errors over longer periods. Hour 3

and hour4 exhibit the highest MDAPE values, indicating challenges in accurately forecasting values over extended hourly periods. The day interval also shows a substantial MDAPE, emphasizing the difficulty of predicting values over a 24-hour period.

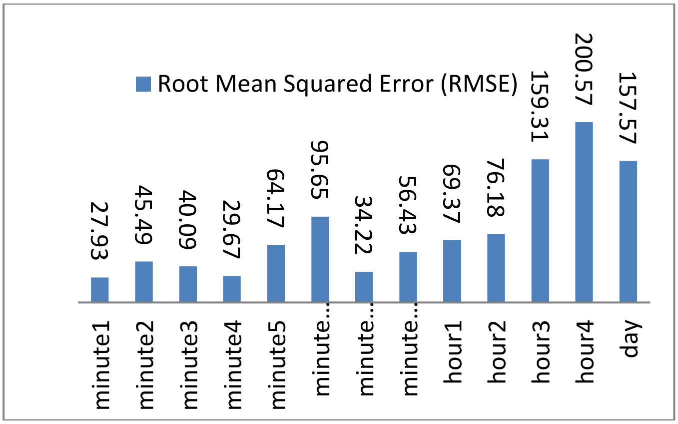

Fig.5. RMSE Values

The observations for Root Mean Squared Error (RMSE) across various time intervals provide insights into the model's forecasting accuracy by measuring the square root of the average squared differences between predicted and actual values is shown in Fig 5. Higher RMSE values indicate larger discrepancies, while lower values suggest better accuracy. In this dataset, the RMSE values for shorter time intervals, such as minutes and hours, generally appear relatively low, indicating reasonably accurate predictions with smaller errors. Figure 6 presents predicted Vs test results.

Fig. 6 Y_Pred vs Y_test

However, as the forecasting horizon extends to a day, the RMSE values notably increase, suggesting a larger dispersion of errors over longer periods. Hour3 and hour4 exhibit the highest RMSE values, indicating challenges in accurately forecasting values over extended hourly periods. The day interval also shows a substantial RMSE, emphasizing the complexity of predicting values over a 24-hour period.

5. Conclusion

In conclusion, the extensive analysis conducted on the provided figures and tables offers a comprehensive insight into the stock's performance and the evaluation metrics of a machine learning model. The stock trend analysis, reveals an overall upward trajectory from February 2015 to February 2024, with periodic instances of volatility, particularly in recent years. The model training and validation, Table 1, reflect a highly effective model with remarkably low training loss (2.78E-07) and a closely aligned validation loss (3.35E-06), emphasizing its accuracy and generalization capabilities. Finally, the comparative analysis provides a visual assessment of the model's predictive accuracy by contrasting predicted and actual prices. Collectively, these findings offer a nuanced understanding of both financial market dynamics and the model's predictive prowess, providing valuable insights for stakeholders and researchers alike.

References

1. W. Lu, J. Li, J. Wang, and L. Qin, ‘A CNN-BiLSTM-AM method for stock price prediction’, Neural Computing and Applications, vol. 33, pp. 4741–4753, 2021.
2. Z. Hu, Y. Zhao, and M. Khushi, ‘A survey of forex and stock price prediction using deep learning’, Applied System Innovation, vol. 4, no. 1, p. 9, 2021.
3. P. Yu and X. Yan, ‘Stock price prediction based on deep neural networks’, Neural Computing and Applications, vol. 32, pp. 1609–1628, 2020.
4. S. Mohan, S. Mullapudi, S. Sammeta, P. Vijayvergia, and D. C. Anastasiu, ‘Stock Price Prediction Using News Sentiment Analysis’, in 2019 IEEE Fifth International Conference on Big Data Computing Service and Applications (BigDataService), 2019, pp. 205–208.
5. M. A. Hossain, R. Karim, R. Thulasiram, N. D. B. Bruce, and Y. Wang, ‘Hybrid Deep Learning Model for Stock Price Prediction’, in 2018 IEEE Symposium Series on Computational Intelligence (SSCI), 2018, pp. 1837–1844.
6. X. Ji, J. Wang, and Z. Yan, ‘A stock price prediction method based on deep learning technology’, International Journal of Crowd Science, vol. 5, no. 1, pp. 55–72, 2021.
7. M. Vijh, D. Chandola, V. A. Tikkiwal, and A. Kumar, ‘Stock closing price prediction using machine learning techniques’, Procedia computer science, vol. 167, pp. 599–606, 2020.
8. S.Aggarwal (2023), LSTM Model Article [Online], Available at:< https://www.projectpro.io/article/lstm-model/832/> , Accessed on Jan 25, 2023.
9. Madhavi, K. Reddy, K. Suneetha, K. Srujan Raju, Padmavathi Kora, Gudavalli Madhavi, and Suresh Kallam. "Detection of COVID 19 using X-ray Images with Fine-tuned Transfer Learning." Journal of Scientific and Industrial Research (2023): 241-248.
10. Polamuri, S.R., Srinnivas, K. & Mohan, A.K. Prediction of stock price growth for novel greedy heuristic optimized multi-instances quantitative (NGHOMQ). *Int J Syst Assur Eng Manag* **14**, 353–366 (2023). https://doi.org/10.1007/s13198-022-01801-3