INDIAN JOURNAL OF SCIENCE AND TECHNOLOGY

ORIGINAL ARTICLE

Predicting Stock Price using Att-Gru-Ols with Vader Sentimental Analysis

Kushagra Mutreja[1], Amiya Ranjan Panda[1]*, Manoj Kumar Mishra[1]

1 School of Computer Engineering, KIIT Deemed to be University, Bhubaneswar, Odisha, India

OPEN ACCESS

Received: 24/01/2025

Accepted: 04/04/2025

Published: 06/07/2025

Citation: Mutreja K, Panda AR, Mishra MK (2025) Predicting Stock Price using Att-Gru-Ols with Vader Sentimental Analysis. Indian Journal of Science and Technology 18(25): 2047-2057. https://doi.org/10.17485/IJST/v18i25.112

* **Corresponding author**.

amiya.pandafcs@kiit.ac.in

Funding: None

Competing Interests: None

Published By Indian Society for Education and Environment (iSee)

ISSN
Print: 0974-6846
Electronic: 0974-5645

Abstract

Objectives: The task of predicting stock prices is more challenging than it seems, especially with the highly volatile markets of today's world. Investors need more robust models for accurate predictions. This paper proposes a novel model that combines sentiment analysis, Ordinary Least Square (OLS), and attention-Gated Recurrent Unit (GRU) to maximise effective prediction. **Methods:** The model integrates a news dataset to calculate sentiment scores and historical stock prices of the BSE Sensex, considering stock value along with sentiment. The GRU captures the crucial non-linear trends in stock movement of the combined dataset. The output of the GRU is passed to an attention mechanism, which selects the most relevant historical input data for predictions. Following this, final stock predictions come from merging outputs of the GRU post-attention mechanism with the original data features and feeding them into an OLS model. **Findings:** Compared with traditional models such as LSTM, Attention BiLSTM, and ARIMA, our model outperformed in metrics for accuracy, such as RMSE, MAE, and R^2, respectively. To capture market sentiment, we examined two approaches: VADER, which is useful for general market sentiment, and FinBERT, which is fine-tuned for financial news. **Novelty:** Real-world tests with livestock data and news revealed additional insights: The model attains better accuracy when sentiment is included, and in the real world, stock prices are governed by many different factors, leading to differences between test dataset accuracy and real-world dataset results.

Keywords: Sentiment Analysis; Attention Mechanism; Stock Price Prediction; Machine Learning; Deep Learning

1 Introduction

Due to the ever-changing nature of the stock market, predicting stock prices accurately remains a major challenge for investors. The Efficient Market Hypothesis (EMH) and Behavioral Finance suggest that stock prices are governed by many underlying factors and market news [1]. Because of this, forecasting stock movements is incredibly difficult due to the market's highly volatile and unpredictable nature [2]. To tackle these challenges, this study leverages machine learning techniques, which include an attention-based GRU model, Ordinary Least Squares (OLS), and sentiment analysis using VADER to enhance prediction accuracy.

Traditionally, stock price forecasting relied on statistical models like ARIMA and GARCH, which are well-suited for handling time-series data [3]. However, as linear models, they struggle to capture the complex, nonlinear patterns present in real-world stock market behavior. This limitation has led researchers to explore machine learning approaches, such as Support Vector Machines (SVM) [4] and Artificial Neural Networks (ANN) [5]. While SVM performs well, its efficiency drops significantly with large datasets due to its computational complexity. ANN, on the other hand, handles larger datasets better but is prone to overfitting and local optimum issues when applied to time-series data.

Long Short-Term Memory (LSTM) networks revolutionized stock market forecasting by capturing long-term dependencies in sequential data. For instance, a 2021 study successfully used an LSTM regression model to predict India's NIFTY 50 index [6]. However, LSTMs come with a high computational cost, making them less practical for large datasets. Meanwhile, Gated Recurrent Units (GRUs), a more efficient alternative, retain much of LSTMs' functionality while being computationally less expensive [7].

In this study, we use a GRU-based sentiment analysis model with an added attention layer to focus on the most relevant historical data points. VADER is applied to extract public sentiment, which is then incorporated into the predictive model to better reflect market dynamics. The GRU's output is combined with OLS regression to capture both short-term fluctuations and long-term trends, allowing for a more comprehensive understanding of stock price behavior.

To place our approach in context, Table 1 provides a summary of recent studies, highlighting the strengths and limitations of different forecasting methods.

Table 1. Summary of recent studies

Study	Method	Advantages	Disadvantages/Research Gap
TOMA LR [3]	ARIMA and GARCH	Effective for handling time series data	Linear assumptions cannot capture non-linear market dynamics
Lu et al. [8]	CNN-BiLSTM-AM	Utilizes attention for improved prediction accuracy	High computational cost; lacks integration of sentiment analysis
Nti et al. [4]	Ensemble Support Vector Machine	Robust classification/regression performance	Scalability issues with large datasets
Jaiswal & Jaiswal [5]	Review on ML techniques	Comprehensive evaluation of various ML models	Highlighted challenges such as overfitting and local optima
Mehtab et al. [6]	LSTM-based deep learning	Captures long-term dependencies effectively	High computational expense; scalability concerns
Chen et al. [7]	Improved GRU-based method	More computationally efficient than LSTM	Lacks integration with sentiment analysis and comprehensive trend modeling

Earlier studies [3]–[5] examined the strengths and weaknesses of traditional statistical methods and the early adoption of machine learning in stock price forecasting. Later research [9]–[10] expanded on this by integrating sentiment analysis and exploring more advanced deep learning models. Building on these efforts, our study compares the proposed approach against existing best practices.

In this paper, the performance of the model is compared with LSTM, standard GRU, ATT-BiLSTM and ARIMA approaches with RMSE, MAE, and R^2 metrics. The main contributions of this study are: Proposing a sentiment based GRU model with attention mechanism to deal with both linear and non-linear modeling The application of VADER and FinBERT for the purpose of stock Sentiment analysis; And the real-life implementation of the model for next day's stock price prediction.

2 Methodology

2.1 Building Model

In this hybrid approach to stock price prediction, the research applied four main techniques: gated recurrent units (GRU), an attention mechanism, ordinary least squares (OLS), and vader sentimental analysis of news-based sentiment. The specific steps involved in model construction are outlined as follows:

2.1.1. Gathering Sentiment Data

We began with collecting the sentiment-based news headlines, from the dataset india-newsheadlines.csv, and used them to calculate the market and economic sentiments relating to various news events.

2.1.2. Sentiment Analysis

Calculated the compound sentiment score using the VADER sentiment analysis tool, creating a quantitative emotional index that complements the stock's trading data.

2.1.3. Preprocess the data

We combined the historical stock data and the sentiment scores to get a multi-source data base. Min-max scalar normalization was performed. We randomly split the data into 70% training and 30% testing for an unbiased evaluation of the model performance.

2.1.4. Initialize parameters

Initialized parameters for GRU such as learning rate, number of iterations, number of hidden layer neurons. The input is fed to the GRU model, which expects to learn complex non-linear patterns. The output is then passed through an attention mechanism that highlights the most salient features in the output of the GRU. The results were then concatenated with historical features and sentiment scores to form a complete dataset.

2.1.5. Integrating OLS

The combined dataset, which includes historical indicators, GRU-attention results, and sentiment scores, is passed to the Ordinary Least Squares (OLS) model. The OLS model provides a linear baseline, complementing the deep learning model's predictions.

2.1.6. Training and Evaluating the Model

The combined model was first evaluated using 5-fold cross-validation to ensure robustness and mitigate overfitting. The dataset was divided into five equal parts, with four folds used for training and one for testing in each iteration. This process was repeated across all folds, and the final performance metrics were computed as the average of all iterations. The cross-validation results yielded an impressive average R^2 score of 0.9883 (± 0.01), along with consistently low RMSE and MAE values, indicating strong performance on historical data. However, to critically assess the real-world applicability of the model, we conducted an additional evaluation using 30 non-consecutive trading days from a different time period that was deliberately excluded from the training and testing phases. Importantly, this real-world test was not intended to demonstrate the robustness of the model**,** but rather to highlight the gap between simulated performance and actual market behaviour. Despite the high R^2 in a controlled setting, real-world testing revealed significant deviations for example, an actual value of 79,043.74 vs. a predicted value of 79,954.91 (Table 5). These discrepancies illustrate that stock prices in real-world scenarios are influenced by complex and often unpredictable factors such as macroeconomic events, policy changes, and investor sentiment shifts that are difficult to capture through historical and sentiment-based features alone. This outcome does not point to overfitting; rather, it exposes the limitations of relying solely on retrospective data for stock price forecasting. The findings stress the importance of cautious interpretation when applying models trained in simulation environments to volatile and multifaceted real-world markets.

2.1.7. Hyperparameter Optimization

To improve the model's performance, we optimized key hyperparameters using grid search based on validation performance. We employed a three-layer GRU model with 64 hidden units per layer, using the tanh activation function to capture sequential dependencies effectively. To prevent overfitting, we incorporated a dropout rate of 0.1 and applied L2 regularization ($\lambda = 0.001$) to the final GRU layer. Additionally, an attention mechanism was introduced, consisting of two dense layers (64 units each) with a tanh activation function, and softmax-based attention scoring to emphasize important time steps. For optimization, we used the Adam optimizer with a learning rate of 0.001, a batch size of 16, and trained the model for 10 epochs. Hyperparameter tuning was performed using Grid Search, systematically evaluating different configurations based on validation performance to identify the best-performing setup. To further refine predictions, we applied Ordinary Least Squares (OLS) regression, integrating GRU outputs with the original feature set.

2.2 Dataset Description

Our model utilizes two decades of historical stock data from the Bombay Stock Exchange (BSE) Sensex, covering the period from January 2, 2001, to December 31, 2020. This dataset includes key trading metrics:

- **Open:** The stock price at the start of the trading day.
- **Close:** The price at market close.

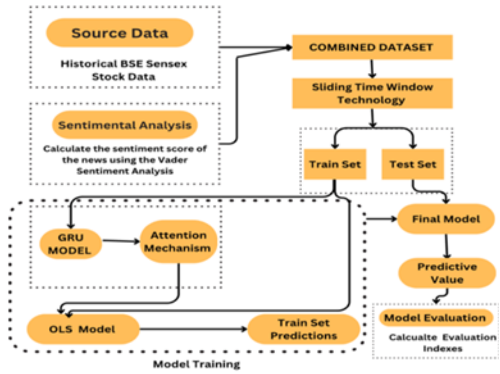

Fig 1. Block Diagram of The Proposed Methodology

- **High:** The highest price reached during the session.
- **Low:** The lowest price recorded during the session.
- **Volume:** The total number of shares traded in a day.

This historical data serves as the foundation of our model, capturing long-term patterns and market volatility over 20 years.

To incorporate market sentiment into our predictions, we use an additional dataset india-news-headlines.csv from Kaggle, which contains daily news headlines related to economic conditions, industry trends, and major market events. Sentiment analysis is performed using VADER (Valence Aware Dictionary and Sentiment Reasoner), which assigns a compound sentiment score ranging from -1 (negative) to +1 (positive). This score acts as a sentiment indicator, allowing the model to assess how public perception might influence stock price movements.

2.3 Machine Learning Models

2.3.1. *The Gated Recurrent Unit (GRU) Model*

GRU, like other recurrent neural networks, is well-suited for processing sequential data such as stock prices. By utilizing update and reset gates, it regulates the flow of information, allowing it to capture long-term dependencies more effectively. This architecture helps address common challenges like the vanishing gradient problem, enabling GRUs to model complex time-series data efficiently[(11)].

2.3.2. *Attention Mechanism*

The Attention mechanism enhances the model's ability to focus on the most relevant data points. Instead of treating all input data equally, it assigns varying levels of importance to different parts of the sequence, ensuring that crucial historical patterns are given more weight. This adaptive approach helps the model prioritize key market events that are more predictive of future stock movements[(12)].

2.3.3. *Ordinary Least Squares (OLS)*

OLS is a classic linear regression technique that finds the best-fitting line by minimizing prediction errors. While GRU captures nonlinear relationships, OLS complements it by modeling linear associations between variables, providing a more comprehensive understanding of stock price trends[(13)].

2.3.4. *Sentiment Analysis Using VADER*

Understanding market sentiment is crucial for predicting stock movements. VADER is a rule-based sentiment analysis tool designed to assess financial news and social media data. By considering context modifiers such as intensity and polarity, VADER delivers an accurate evaluation of public sentiment, helping to capture investor emotions and market trends[(14)].

3 Results and Discussion

3.1 Evaluation Metrics

3.1.1. Root Mean Square Error (RMSE)

RMSE is a common measure of the difference between values predicted by the model and the actual observed values. It emphasizes larger errors due to its squaring function, making it useful for understanding the scale of errors in predictions.

$$\text{The equation for RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$$

3.1.2. Mean Absolute Error (MAE)

MAE provides the average absolute difference between predicted values and observed values, giving a straightforward measure of prediction accuracy without emphasizing large errors as RMSE does.

$$\text{The equation for MAE} = \frac{1}{n}\sum_{i=1}^{n} y_i - \hat{y}_i \vee$$

3.1.3. Coefficient of Determination (R^2)

The R^2 metric represents the proportion of the variance in the dependent variable that is predictable from the independent variables. It ranges from 0 to 1, with higher values indicating better model performance in capturing the variability in the data.

$$\text{The equation for } R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y_i})^2}{\sum_{i=1}^{n}(y_i - \overline{y}_i)^2}$$

where:

- y_i is the actual observed value,
- $\hat{y}_i$ is the predicted value,
- n is the number of observations.
- $\overline{y}_i$ is the mean of observed values.

3.2 Results of the experiment

3.2.1. Comparison with other models

Several studies have highlighted the effectiveness of sentiment analysis and deep learning models in stock price prediction. Heiden et al.[9] demonstrated how integrating LSTM with VADER sentiment analysis enables models to capture temporal patterns driven by market sentiment. Similarly, Akhila Mamillapalli et al.[15] introduced a GRU-based model with VADER, emphasizing the importance of combining sentiment data with GRU to enhance predictive accuracy. While these approaches showcase the potential of sentiment-based models, our findings suggest that integrating OLS with attention mechanisms further improves prediction performance.

Han and Fu[16] leveraged an Attention BiLSTM model to forecast stock prices, effectively capturing contextual relationships using a bidirectional structure. However, their approach struggled to prioritize the most influential factors driving stock price fluctuations. Our Attention-GRU model addresses this limitation by incorporating VADER sentiment analysis and OLS, leading to better predictive accuracy.

Minhaj et al.[10] explored the combination of sentiment analysis with traditional statistical models, specifically ARIMA, for stock price forecasting. While their approach demonstrated the usefulness of sentiment data in statistical models, deep learning methods such as GRU and Attention-GRU are better equipped to model complex temporal patterns, outperforming traditional techniques.

We compared our proposed Attention-based GRU model with Ordinary Least Squares (OLS) regression and VADER sentiment analysis to several other popular deep learning models that also used VADER sentiment analysis on news data, as indicated in Table 2. The comparison was based on three evaluation metrics: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and the Coefficient of Determination (R^2), each capturing different dimensions of model performance.

Table 2. Comparison of Models

Sl. No.	Model Name
1.	LSTM with VADER
2.	GRU with VADER sentiment analysis
3.	CNN - ATT BiLSTM
4.	ARIMA

Figures 2 to 6 illustrate the comparison between the predicted and actual stock prices bydifferent models. These figures show the number of trading days in the test set on the X-axis and the closing price of the stock on the Y-axis. In these figures, the blue line indicates the predicted closing price, the orange line represents the actual closing price for the baseline models, and the red line denotes the actual closing price in the proposed model.

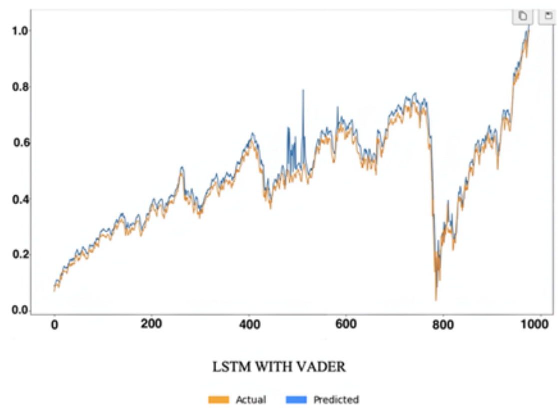

Fig 2. Predicted Value Of BSE-Sensex Using LSTM And Vader (Blue: Predicted, Orange: Actual)

Fig 3. Predicted Value Of BSE-Sensex Using ATT-BiLSTM (Blue: Predicted, Orange: Actual)

Table 3 shows the experimental results of the Attention-based GRU Model in combination with Vader and OLS. We can see that our proposed model has the highest RMSE, MAE and R2 Score values.

This superior result emphasizes the originality of our method. By combining the non-linear modeling skills of an attention-enhanced GRU with the linear trend-capturing strength of OLS regression and the contextual insights supplied by VADER sentiment analysis, our approach fills a significant gap in the existing literature. While previous research has primarily focused on deep learning techniques or traditional statistical methods with sentiment enhancements, our integrated approach captures complex temporal dependencies while also effectively incorporating market sentiment, resulting in more robust and reliable stock price forecasting.

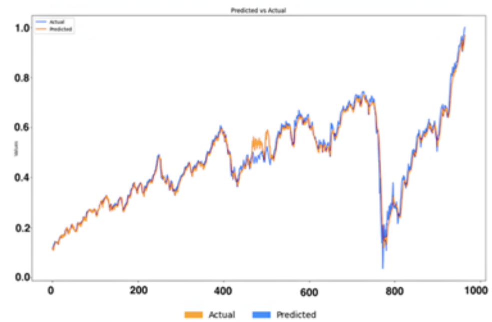

Fig 4. Predicted Value Of BSE-Sensex Using ARIMA (Blue: Predicted, Orange: Actual)

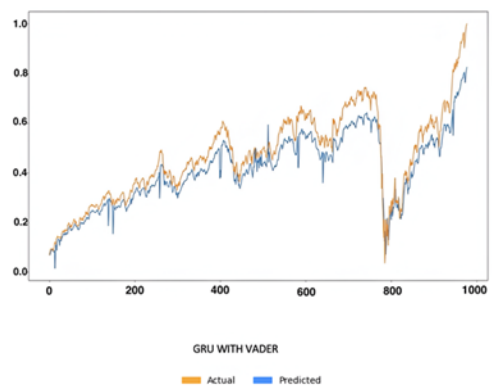

Fig 5. Predicted Value Of BSE-Sensex Using GRU And Vader (Blue: Predicted, Orange: Actual)

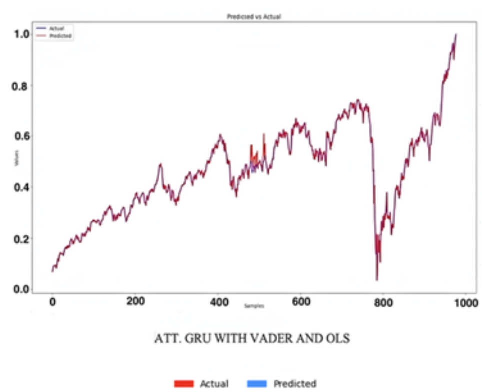

Fig 6. Predicted Value Of BSE-Sensex Using ATT-GRU WITH OLS And Vader (Blue: Predicted, Red: Actual)

Table 3. Performance metrics for different models

Model Name	RMSE	MAE	R^2
LSTM With Vader [9]	0.0369	0.0282	0.9571
GRU with Vader [15]	0.0682	0.0571	0.8535
ATT. BILSTM [16]	0.1104	0.1038	0.6163
ARIMA [10]	0.0258	0.0172	0.9779
ATT. GRU with Vader and OLS (Proposed)	0.0194	0.0124	0.9881

In conclusion, the innovations provided in our study specifically, the combination of attention mechanisms, sentiment analysis, and OLS regression with a GRU model show significant gains over previous efforts. This comprehensive framework considerably improves the accuracy of next-day stock price projections, establishing a new standard in the industry.

3.2.2. Comparison with Attention GRU Model Using FinBERT Sentiment Analysis and OLS

To validate the effectiveness of our Attention-GRU model with VADER sentiment analysis and OLS, we compared it against an alternative approach that replaced VADER with FinBERT for sentiment analysis. Given that FinBERT is specifically designed for financial contexts, it was expected to provide a more precise sentiment representation of market-related text. While the FinBERT-based model delivered strong results with an RMSE of 0.0205, MAE of 0.0127, and R^2 of 0.9867**,** our VADER-based model outperformed it slightly with an RMSE of 0.0194, MAE of 0.0124 and R^2 of 0.9880. Several key factors contribute to this marginal improvement.

First, VADER has demonstrated robustness across various text domains, making it particularly effective in capturing sentiment trends in stock-related news, where sentiment shifts are not always domain-specific. Unlike FinBERT, which is fine-tuned for financial texts, VADER performs well across diverse sources, including general news headlines that significantly impact stock prices. Additionally, VADER's lexicon-based approach efficiently captures sentiment polarity and intensity, aligning well with broad financial news rather than being limited to technical financial documents. Since stock price movements are influenced by a mix of financial reports and broader market events, VADER's adaptability in processing a variety of news sources may have contributed to the improved predictive accuracy of our model.

This comparison highlights that while FinBERT excels in financial contexts**,** VADER's simplicity and versatility make it a compelling choice for stock price forecasting models that incorporate wider news sentiment to enhance prediction performance.

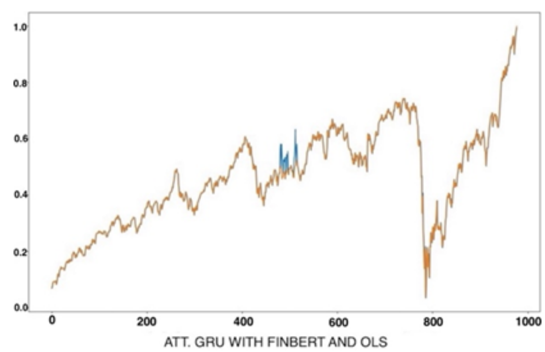

Fig 7. Predicted Value of BSE-Sensex using ATT-GRU with OLS and FINBERT

3.2.3. Real-world Testing

To assess the practical applicability and real-world limitations of the proposed model, we conducted testing using stock market data and corresponding news sentiment for 30 non-consecutive trading days (Nov 28, 2024 to Oct 2024)**.** This evaluation was designed not to further validate the model's robustness but to investigate how a model that performs well in a controlled environment handles the uncertainty and complexity of real-world markets**.**

A web-based application was developed to facilitate real-world testing, which serves as a user-friendly interface for testing the model. The trained model was saved using Python's Pickle module and integrated into the web application for seamless deployment. Users can input relevant stock data and news sentiments into the web app, which then uses the pickled model to predict stock closing prices in real time. This setup allowed us to evaluate the model's performance in an environment closely resembling its intended real-world application.

The error bar graph visualizes the model's predicted stock closing prices compared to the actual closing prices over a 30-day period. While the model demonstrated high accuracy during cross-validation, the graph reveals significant deviations in real-world predictions. Each error bar represents the absolute difference between the predicted and actual prices for a given date. The inconsistencies emphasize that despite strong performance in a simulated environment, stock price prediction models may underperform in real-world settings due to unpredictable market influences such as economic events, investor behavior, and geopolitical developments.

Table 4. Stock Data for Selected Dates (First 10 of 30 Records)

DATE	Open Price	High	Low	Closed price	Volume
Nov 28,2024	80,281.64	80,447.40	78,918.92	79,043.74	7,900
Nov 27,2024	80,121.03	80,511.15	79,844.49	80,234.08	5,700
Nov 26,2024	80,415.47	80,482.36	79,798.67	80,004.06	10,400
Nov 25,2024	80,193.47	80,473.08	79,765.99	80,109.85	14,000
Nov 22,2024	77,349.74	79218.19	77,226.69	79,117.11	17,100
Nov 21,2024	77,711.11	77,711.11	76,802.73	77,155.79	8,800
Nov 19,2024	77,548.00	78,451.65	77,411.31	77,578.38	7,200
Nov 18,2024	77,863.54	77,886.97	76,965.06	77,339.01	8,200
Nov 14,2024	77,636.94	78,055.52	77,424.81	77,580.31	10,100
Nov 13,2024	78,495.53	78,690.02	77,533.30	77,690.95	7,400

Table 5. Comparison of Actual and Predicted Closing Prices (First 10 of 30 Records)

DATE	Actual Closed price	Predicted closing Price
Nov 28,2024	79,043.74	79954.91
Nov 27,2024	80,234.08	79659.93
Nov 26,2024	80,004.06	79882.16
Nov 25,2024	80,109.85	79226.24
Nov 22,2024	79,117.11	76898.32
Nov 21,2024	77,155.79	77376.05
Nov 19,2024	77,578.38	77059.51
Nov 18,2024	77,339.01	77405.92
Nov 14,2024	77,580.31	77412.34
Nov 13,2024	77,690.95	78321.67

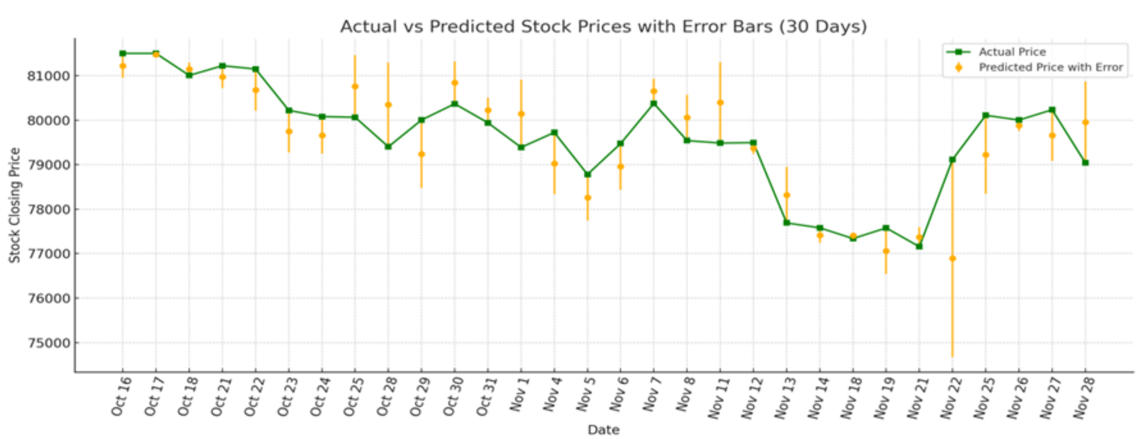

Fig 8. Actual vs Predicted Stock Prices with Error Bars (30 Days)

3.3 Discussion

3.3.1. Comparing Our Model with Existing Approaches

The results show that our attention-based GRU with VADER and OLS model outperforms previous approaches in stock price prediction. With the lowest RMSE (0.0194) and MAE (0.0124), as well as the highest R2 score (0.9881), it demonstrates strong predictive accuracy and a significant reduction in error compared to existing methods.

When compared to LSTM with VADER[9], which had an RMSE of 0.0369 and an R2 of 0.9571, our model shows notable improvements. The integration of attention mechanisms likely played a key role in enhancing predictive power by helping the model focus on relevant data points while filtering out less useful information.

The GRU with VADER model[15] showed lower performance, with an RMSE of 0.068**2** and an R2 of 0.8535. While GRUs are known for being efficient, they might not always capture long-term dependencies as effectively as LSTMs, especially without attention mechanisms to guide them.

Interestingly, the Attention-based BiLSTM model[16] had the weakest performance, with an RMSE of 0.1104 and an R2 score of 0.6163. While bidirectional LSTMs can capture patterns in both directions, their added complexity may have led to overfitting rather than improving predictions in this particular task.

One surprising takeaway from our results is the strong performance of the ARIMA model[10], which had an RMSE of 0.0258 and an R2 of 0.9779. Despite being a traditional statistical model, ARIMA still competes well with deep learning models in short-term forecasting. However, its reliance on linear relationships means it may struggle with the nonlinear dynamics often seen in financial markets, where deep learning models have an edge.

3.3.2. The Role of Sentiment Analysis and Model Design

Our study confirms that sentiment analysis enhances stock price prediction by capturing market sentiment and investor psychology. While FinBERT performed well (RMSE: 0.0205, MAE: 0.0127, R^2: 0.9867), our VADER-based model achieved slightly better results (RMSE: 0.0194, MAE: 0.0124, R^2: 0.9880).

This improvement can be attributed to VADER's efficiency in handling short financial texts such as news headlines and social media posts, which often drive market movements. Additionally, VADER is computationally lightweight, enabling faster processing compared to FinBERT's transformer-based architecture.

Moreover, our model's attention mechanism prioritized key historical patterns, while the OLS adjustment step fine-tuned predictions, further reducing errors. These results highlight that a hybrid approach, combining deep learning with statistical techniques, leads to more accurate predictions.

3.3.3. ***Implications for Real-World Use***

The error bar analysis offers a visual narrative of the gap between predicted and actual stock prices across a 30-day real-world testing period. Despite the model's high R^2 score and consistent performance during cross-validation, the variation in prediction accuracy observed here reinforces the limitations of relying solely on historical and sentiment data. While some predictions were closely aligned with actual outcomes, others showed considerable deviations, indicating that the market's complexity cannot be fully captured by past trends and sentiment indicators alone. The presence of external forces ranging from macroeconomic policy changes and geopolitical developments to investor psychology and market speculation introduces volatility that defies traditional modeling techniques. These findings highlight a crucial consideration for deploying such models in practical applications: they should be used as decision-support tools rather than decision-making authorities. Incorporating real-time data sources, adaptive learning mechanisms, or even hybrid approaches that blend quantitative and qualitative insights may enhance the robustness and reliability of such predictive systems in live market conditions.

4 Conclusion and Future Work

In this study, we introduced a novel stock price prediction framework that integrates traditional financial indicators and sentiment analysis, evaluating its effectiveness across three key dimensions: real-world testing, sentiment analysis comparison, and model performance analysis. Our proposed model demonstrated superior predictive accuracy compared to several baseline approaches. By effectively capturing non-linear market behaviors and incorporating market sentiment, our model achieved an RMSE of 0.0194, MAE of 0.0124, and R^2 of 0.9881, significantly surpassing conventional models like ARIMA (RMSE = 0.0258, MAE = 0.0172, R^2 = 0.9779) and deep learning models such as LSTM (RMSE = 0.0369, MAE = 0.0282, R^2 = 0.9571).

The combination of GRU, OLS, an attention mechanism, and sentiment analysis has proven to be a robust framework for stock price forecasting. A comparative analysis of sentiment analysis methods revealed that VADER outperformed FinBERT in capturing public sentiment toward stock price movements, making it a more effective sentiment analysis tool for this application.

During real-world testing, the model accurately predicted closing prices on certain trading days, reinforcing its practical utility. However, significant deviations were observed on other days, highlighting the challenges posed by the volatile and unpredictable nature of financial markets. These discrepancies emphasize the influence of external factors, such as macroeconomic trends, geopolitical developments, and unexpected market shifts, which extend beyond the scope of sentiment analysis and historical stock data.

To further enhance the robustness of stock price prediction models, future research should examine the role of external factors such as global events, industry trends, and regulatory changes in shaping market behavior. A dynamic model capable of adapting to sudden market fluctuations would improve real-world applicability. In this regard, reinforcement learning techniques could be explored as a potential solution to create adaptive and resilient financial prediction models in highly volatile trading environments.

References

1) Gu G. The Dynamic Interplay of Market Forces and Human Behavior: A Critical Review of Efficient Market Hypothesis and Behavioral Finance. *Advances in economics, management and political sciences*. 2023 Nov;31(1):55–60. https://doi.org/10.54254/2754-1169/31/20231498.
2) Sawhney R, Agarwal S, Wadhwa A, Shah RR. Deep Attentive Learning for Stock Movement Prediction From Social Media Text and Company Correlations. 2020;p. 8415–8426. Available from: https://aclanthology.org/2020.emnlp-main.676/.
3) TOMA LR. Exploring the Effectiveness of ARIMA and GARCH Models in Stock Price Forecasting: An Application in the IT Industry. *Informatica Economica*. 2023 Sep;27(3/2023):61–72. Available from: https://revistaie.ase.ro/content/107/05%20-%20toma.pdf.
4) Nti IK, Adekoya AF, Weyori BA. Efficient Stock-Market Prediction Using Ensemble Support Vector Machine. *Open Computer Science*. 2020 Jul;10(1):153–163. https://doi.org/10.1515/comp-2020-0199.
5) Jaiswal S, Jaiswal T. Review on Machine Learning Techniques for Stock-Market Forecasting. *Artificial Intelligence Evolution*. 2020 May;p. 34–47. https://doi.org/10.37256/aie.112020317.
6) Mehtab S, Sen J, Dutta A. Stock price prediction using machine learning and LSTM-based deep learning models. Berlin, Germany. Springer. 2021. https://doi.org/10.1007/978-981-16-0419-5_8.
7) Chen C, Xue L, Xing W. Research on Improved GRU-based Stock Price Prediction Method. 2023. https://doi.org/10.20944/preprints202307.0472.v1.
8) Lu W, Li J, Wang J, Qin L. A CNN-BiLSTM-AM method for stock price prediction. *Neural Computing and Applications*. 2020 Nov;33(10):4741–4753. https://doi.org/10.1007/s00521-020-05532-z.
9) Heiden A, Parpinelli RS. Applying LSTM for Stock Price Prediction with Sentiment Analysis. 2021. https://doi.org/10.21528/cbic2021-45.
10) Minhaj N, Ahmed R, Khalique IA, Imran M. A Comparative Research of Stock Price Prediction of Selected Stock Indexes and the Stock Market by Using Arima Model. *Global Economics Science*. 2022 Oct;p. 1–19. https://doi.org/10.37256/ges.4120231426.
11) Dey P, Hossain E, Hossain M, Chowdhury MA, Alam M, Hossain MS, et al. Comparative Analysis of Recurrent Neural Networks in Stock Price Prediction for Different Frequency Domains. *Algorithms*. 2021 Aug;14(8):251. https://doi.org/10.3390/a14080251.
12) Kelvin K, Sinaga FM, Winardi S, Susmanto S. Exploring New Frontiers: XCEEMDAN, Bidirectional LSTM, Attention Mechanism, and Spline in Stock Price Forecasting. *Indonesian Journal of Artificial Intelligence and Data Mining*. 2024 Jun;7(2):384. https://doi.org/10.24014/ijaidm.v7i2.29649.
13) Li Y. Stock Price Prediction based on Multiple Regression Models. *Highlights in Science, Engineering and Technology*. 2023 Apr;39:657–662. https://doi.org/10.54097/hset.v39i.6622.
14) Cristescu MP, Nerișanu RA, Mara DA. Using Data Mining in the Sentiment Analysis Process on the Financial Market. *Journal of Social and Economic Statistics*. 2022 Dec;11(1-2):36–58. https://doi.org/10.2478/jses-2022-0003.
15) Mamillapalli A, Ogunleye B, Inacio ST, Shobayo O. GRUvader: Sentiment-Informed Stock Market Prediction. *Mathematics*. 2024 Nov;12(23):3801. https://doi.org/10.3390/math12233801.
16) Han C, Fu X. Challenge and Opportunity: Deep Learning-Based Stock Price Prediction by Using Bi-Directional LSTM Model. *Frontiers in Business, Economics and Management*. 2023 Apr;8(2):51–54. https://doi.org/10.54097/fbem.v8i2.6616.