Available Online at www.ijcsmc.com

International Journal of Computer Science and Mobile Computing

A Monthly Journal of Computer Science and Information Technology

ISSN 2320–088X
IMPACT FACTOR: 7.056

IJCSMC, Vol. 13, Issue. 10, October 2024, pg.32 – 44

A Novel Deep Learning Model for Stock Prediction Integrating Historical Prices, Core Volatility Trends and Financial News

Anshuman Tekriwal

Independent Researcher, India
tekriwalanshuman@gmail.com

DOI: https://doi.org/10.47760/ijcsmc.2024.v13i10.004

Abstract*— Accurate stock price forecasting is crucial for investors and analysts aiming to make informed decisions and mitigate risks in financial markets. While deep learning models have proven effective in identifying complex patterns in quantitative data, they often struggle to account for non-quantitative factors such as corporate events, market sentiment, and regulatory changes that can cause unpredictable stock price movements. To address this challenge, an ensemble model is introduced that integrates historical price data, financial news, and stock fundamentals to predict next-day stock prices. The model combines one-dimensional Convolutional Neural Networks (1D CNNs), Long Short-Term Memory networks (LSTMs), and Gated Recurrent Units (GRUs), applied independently to each data stream and concurrently as a unified ensemble, to capture both short-term and long-term market behaviors. NVIDIA (NVDA) stock was selected for testing due to its high volatility and rapid price fluctuations, presenting a challenging case for predictive modeling. Using data spanning from April 2014 to March 2024, the model achieved a coefficient of determination (R^2) of 0.983 and a mean absolute error (MAE) of $12.72. These results suggest that the ensemble approach effectively captures both quantitative and qualitative factors, making it a promising tool for improving stock price prediction in volatile market conditions.

***Keywords*— Stock price prediction, Deep learning, Financial news analysis, Ensemble model, Time series forecasting**

I. INTRODUCTION

Accurate prediction of stock prices is a complex challenge due to the multitude of factors influencing market dynamics. Fundamentally, stock prices are determined by the forces of supply and demand: higher demand for a stock increases its price, while lower demand decreases it. Fluctuations in stock demand are affected by numerous factors, including company-related news such as leadership changes, new product announcements, and innovations; sector-wide trends like regulatory changes, industry booms, and demographic shifts; market trends including investor sentiment; economic indicators like inflation rates and interest rate changes; and unexpected events such as geopolitical tensions and natural disasters. In addition to these qualitative factors, stock fundamentals—such as financial performance (revenue, cash flow, profitability), financial health (borrowing rates, debt-to-equity ratios, liquidity ratios), growth potential, and stock volatility—play a crucial role in determining the frequency and magnitude of stock price movements. Therefore, gathering and processing this multifaceted information to make accurate predictions presents a significant challenge.

Recent advancements in Natural Language Processing (NLP) offer opportunities to leverage financial news as an effective source for capturing potential changes in stock prices due to qualitative indicators. The development of network architectures such as Long Short-Term Memory networks (LSTMs) and Gated Recurrent Units (GRUs), along with powerful embedding and encoding techniques, enables models to capture general sentiment towards a stock through news articles, thereby enhancing stock price predictions. However, existing models often fail to predict stock data accurately because they focus on a single type of data or combine only textual and historical information, rather than integrating stock news, historical prices, and stock fundamentals comprehensively. Considering textual or numerical information separately may not provide the model with all relevant data, potentially resulting in less reliable predictions.

This paper proposes an ensemble model that integrates one-dimensional Convolutional Neural Networks (1D CNNs), bidirectional LSTMs, and GRUs to process stock fundamentals, financial news, and historical prices. By combining these diverse data sources, the proposed ensemble architecture aims to maximize the capture of stock and market behavior, improving prediction accuracy.

To the best of our knowledge, no prior research has utilized a blend of historical prices, technical and fundamental indicators, and financial news to predict stock price movements in this manner.

The contributions of this paper are as follows:

- Assimilates historical, analytical, and textual financial data to capture a comprehensive view of market conditions.
- Develops an ensemble model that leverages state-of-the-art neural network architectures to process a diverse dataset.
- Addresses existing limitations in understanding the effect of qualitative data on stock prices through the proposed innovations.

The remainder of this paper is organized as follows: Section 2 provides an overview of related work; Section 3 outlines the methodology involved in this research; Section 4 presents the experimental setup and corresponding results; Section 5 offers insights on the findings and discussions; and Section 6 concludes the paper with a summary of the findings and their implications.

II. RELATED WORK

A. *Conventional Approaches to Time Series Forecasting*

Traditional methods for time series forecasting primarily rely on statistical models that process historical data to predict future values. Techniques such as AutoRegressive Integrated Moving Average (ARIMA) (Ariyo et al., 2014) and Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) models are commonly employed. ARIMA captures the relationships between observations and their lagged values, while GARCH models model time-varying volatility. Although effective for stationary data and volatility modeling, these methods have significant limitations. They assume linear relationships, struggle with non-linear dependencies, and are inadequate for sudden market shifts, rendering them less suitable for the complexities of financial markets (Zhang, 2021).

B. *Deep Learning Approaches*

The advent of deep learning models has significantly improved time series forecasting methods due to their ability to capture complex patterns and dependencies in data. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory networks (LSTMs) and Gated Recurrent Units (GRUs), have transformed this field by effectively modelling sequential data. LSTMs are an evolution of RNNs designed to overcome the vanishing and exploding gradient problems inherent in traditional RNNs (Fischer & Krauss, 2018). By utilizing memory cells, LSTMs can capture both short-term and long-term dependencies within sequential data. Consequently, they have been extensively used in financial forecasting despite their computational intensity (Siami-Namini & Siami Namin, 2018). GRUs are a variant of the LSTM architecture that simplifies the network by combining the forget and input gates into a single update gate. This results in a more computationally efficient model while retaining the ability to capture long-term dependencies. GRUs have demonstrated comparable performance to LSTMs in many applications, including stock price prediction. Their simpler structure allows for faster training and reduced computational overhead, making them an attractive alternative to LSTMs (Chen et al., 2015). However, while deep learning models excel at modelling complex temporal dependencies, they often rely solely on historical price data, potentially overlooking other influential factors such as financial news and stock fundamentals.

C. *Ensemble Models*

Ensemble modeling has gained prominence in stock price prediction due to its ability to enhance accuracy and robustness by combining multiple learning algorithms (Pasupulety et al., 2021). Techniques like bagging and boosting improve prediction performance by reducing variance and addressing bias, focusing on correcting previous errors. Recent studies have employed ensemble methods to integrate different models or data sources, but few have comprehensively combined diverse data types such as historical prices, technical indicators, and textual information.

D. *Data Integration in Stock Price Forecasting*

Accurate prediction of stock prices requires the integration of diverse data sources to capture a comprehensive view of market conditions. These sources include historical stock prices, stock fundamentals and technical indicators, as well as financial news and sentiment analysis. Historical stock prices provide the foundation for time series analysis and forecasting models. They include daily closing prices, high and low prices, trading volumes, and adjusted prices for dividends and splits. Historical data is crucial for identifying trends, patterns, and seasonality in stock movements (Fama, 1970). While essential, relying solely on historical data can be insufficient, as it may not account for current conditions or sudden market shocks (Campbell et al., 1997). Stock fundamentals encompass a company's financial statements, including balance sheets, income statements, and cash flow statements. Key metrics such as earnings per share (EPS), price-to-earnings (P/E) ratio, and return on equity (ROE) provide insights into a company's financial health and performance (Lev & Thiagarajan, 1993). Technical indicators, derived from historical price and volume data, such as moving averages, Relative Strength Index (RSI), and Moving Average Convergence Divergence (MACD), help identify momentum, overbought or oversold conditions, and potential reversal points (Murphy, 1999). These indicators offer a current view of the conditions specific to the concerned stock. Financial news and sentiment analysis capture the market's reaction to events and announcements that can impact stock prices. This includes news articles, press releases, earnings reports, and social media sentiment. By analyzing the tone and context of these texts, models can incorporate market sentiment into their predictions (Tetlock, 2007). Advances in Natural Language Processing (NLP) algorithms provide tools to understand the qualitative factors behind market dynamics (Bollen et al., 2011).

E. *Gaps in Existing Research*

While numerous studies have utilized individual data sources for stock price prediction, few have integrated multiple types of data comprehensively. Existing models often focus on either quantitative data, such as historical prices and technical indicators, or qualitative data, like financial news and sentiment analysis. For instance, Hu et al. (2018) employed deep learning models with sentiment analysis to predict stock movements but did not integrate stock fundamentals. Similarly, Zhang et al. (2020) used LSTMs on historical price data but did not consider textual information. Moreover, while ensemble models have been used to combine different algorithms, there is a lack of research on ensemble architectures that integrate diverse data streams using state-of-the-art neural network models tailored to each data type. This gap highlights the need for a holistic approach that assimilates historical prices, technical and fundamental indicators, and financial news to improve prediction accuracy.

III. METHODOLOGY

This section provides a general overview of the methodology involved in our research, and offer insight on the NLP tools and techniques used in data ingestion and preprocessing. Furthermore, it highlights the architectures used in defining the model, the process of training it as well as model evaluation.

A. *Data Collection*

Metrics related to the stock of NVIDIA Corporation (NASDAQ: NVDA) from April 1, 2014, to March 1, 2024—approximately ten years of data—were used for this research.

1) *Historical Data:* The historical market data for NVDA was obtained from the yfinance library (Yahoo! Finance API). The data points included daily open, high, low, and close prices, as well as trading volume. This data allows the model to learn from previous stock behavior and to identify recurring patterns and seasonality in the stock's movements.
2) *Stock Fundamentals and Technical Indicators:* The stock fundamentals for NVDA were sourced from the ORATS Expanded Core Option Volatility Data Feed (OPT2), which contains data on implied and historical volatility readings. These data points include several crucial measures such as historical and implied volatility, earnings move studies, percentile analyses, and borrow rates. Incorporating these indicators enhances the model's ability to accurately forecast stock prices by providing a comprehensive view of past and present market conditions and potential outcomes.
3) *Financial News:* Financial news related to NVDA was sourced from the EODHD Financial News API and scraped from the NVIDIA Newsroom Archive. These data sources provided extensive coverage of financial news articles, press releases, and social media sentiment. This raw data provides general market sentiment

toward NVDA stock, as well as information on qualitative factors such as innovations or scandals that may affect the company's stock price. Thus, financial news plays a pivotal role in supplying the model with information that may result in sudden price fluctuations.

B. *Data Preprocessing*

1) *Historical Data:*

Handling Missing Values: Missing values due to weekends and market closure days were replaced using the backward fill method. This ensured continuity in the dataset by filling missing entries with the most recent available data points.

Normalization: The Min-Max scaling technique was applied to scale all market prices between 0 and 1. This reduction in scale minimizes the effect of large price variations over the years and facilitates faster convergence during model training.

2) *Stock Fundamentals & Technicals:*

Handling Missing Values: Similar to the historical data, missing values were addressed using the backward fill method to maintain dataset continuity.

Normalization: Min-Max scaling was utilized to scale all indicators between 0 and 1. This ensures that no single feature disproportionately influences the model's distance calculations and makes the data easier to interpret.

Feature Selection: A correlation matrix with respect to the target price was computed, and the 20 most correlated indicators were selected for use. Table 1 lists the technical and fundamental indicators employed in the model.

The correlation matrix to the target price was computed and the 20 most closely related indicators were selected for use.

TABLE I
TECHNICAL INDICATORS USED

Name	Description
cvolu	Total call open interest
pvolu	Put options volume for all strikes
slope	Best-fit regression line through the strike volatilities adjusted to the tangent slope at the 50 delta. The slope is the change in the implied volatility for every 10 delta increase in the call delta
deriv	Derivative or curvature of the monthly strikes at 30 day interpolated. The derivative is the change in the slope for every 10 delta increase in the call delta
confidence	Weighted confidence derived from number of options and bid-ask width for each expiration
impliedlearningsmov	Implied move after earnings announcement based on the straddle price less the residual straddle price estimate for after earnings
pxatmiv	Average bid-ask price near the close
clspx1m	Price at the prior month (21 trading days ago)
stkpxchng1m	Percentage of price change over the prior month (21 trading days ago)
clshvxern10d	10-day historical close-to-close volatility ignoring earnings moves
clshvxern60d	60-day historical close-to-close volatility ignoring earnings moves
clshvxern120d	120-day historical close-to-close volatility ignoring earnings moves
clshvxern252d	252-day historical close-to-close volatility ignoring earnings moves
contango	Slope of the short-term at-the-money implied volatilities ex-earnings
iv90d	90 calendar-day interpolated implied volatility at the 50 delta
iv1y	365 calendar-day interpolated implied volatility at the 50 delta
ffexern90_60	Flat forward ex-earnings volatility extracted from the 60-day and 90-day implied ex-earnings volatility
ffexern180_90	Flat forward ex-earnings volatility extracted from the 180-day and 90-day implied ex-earnings volatility
fbfwd90_30	Flat forward volatility divided by the forward volatility both extracted from the 30-day and 90-day implied volatility

3) *Financial News:*

Text Cleaning: The textual data was cleaned by removing stop words, lemmatizing words to their base forms, stemming, and eliminating punctuation. This preprocessing reduces noise and standardizes the text for effective analysis.

Tokenization and embedding: The cleaned text was tokenized using TensorFlow's tokenizer and converted into numerical vectors. GloVe pre-trained word embeddings were employed to represent words in a continuous vector space, capturing semantic relationships.

4) *Final Data Preparation:*

All data types were concatenated and rolling windows of size 25 were created to enable the model to detect underlying patterns in price changes effectively. The final processed dataset spanned from April 1, 2014, to March 1, 2024, containing a total of 2,824 observations.

TABLE II
SUMMARY OF EACH DATA TYPE IN THE DATASET

Name	Data Shape	Values Range
Historical Prices	(2798, 25, 1)	0 to 1
Stock Fundamentals	(2798, 25, 20)	0 to 1
Financial News	(2798, 1000)	0 to 1000
Target '*y*'	(2798, 1, 1)	0 to 1

C. *Model Architecture*

The model predicting the stock price of NVIDIA Corporation (NVDA) is designed to capture both temporal and qualitative patterns in the data. This section provides a detailed overview of each layer in the neural network and the rationale behind their selection.

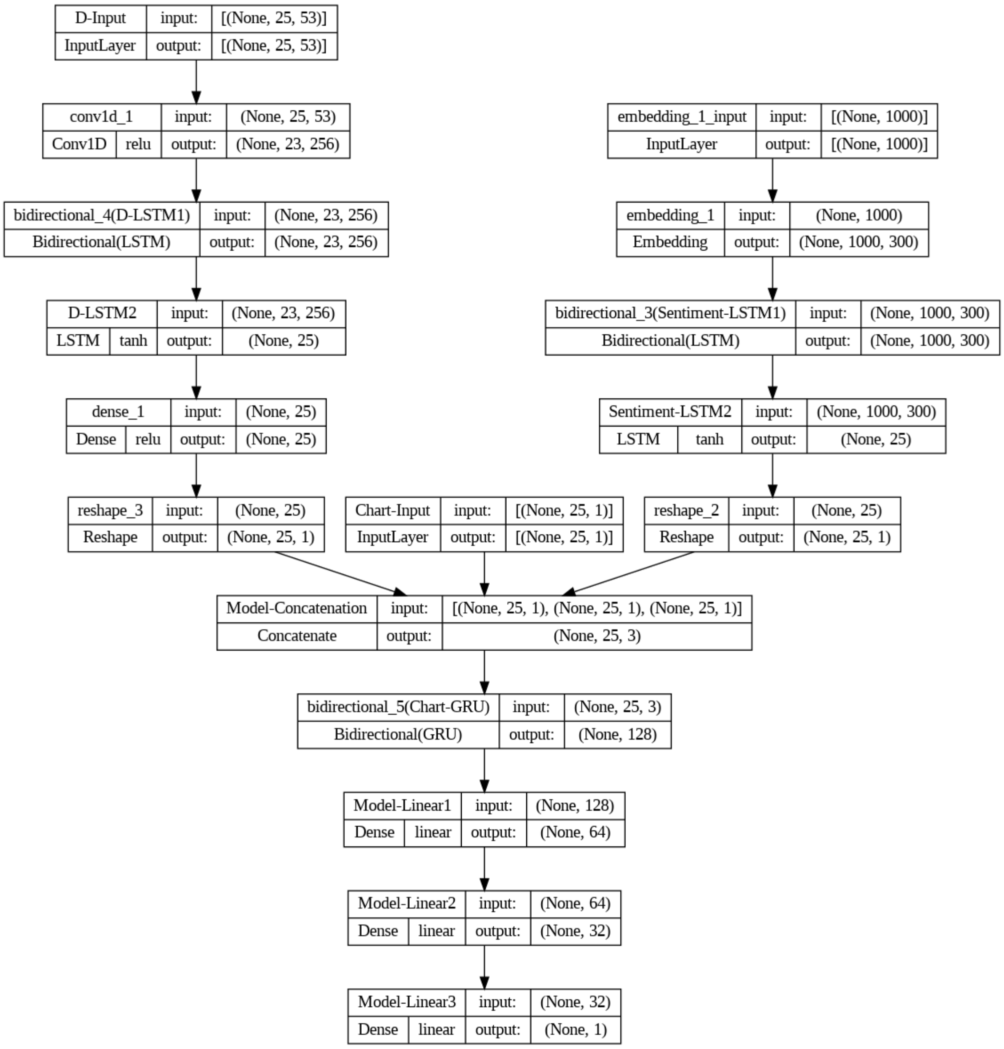

Fig. 1 Model Architecture

The input layer for stock fundamental data, referred to as D-Input, consists of 53 features per timestamp. This data is first processed through a one-dimensional Convolutional Neural Network (Conv1D) layer, which enables the model to learn local patterns and trends over short time windows. The output from the Conv1D layer is then passed into a Bidirectional Long Short-Term Memory (BiLSTM) layer, capturing temporal dependencies from both past and future time steps. This ensures the model considers context in both directions. The BiLSTM further refines the temporal features extracted. A subsequent Dense layer adds non-linearity to the model and helps mitigate the vanishing gradient problem. Finally, the output is reshaped to ensure compatibility for concatenation with weights from other data streams.

The Chart-Input layer handles the input for historical closing prices of NVDA stock, consisting of the previous 25 time steps. This provides a comprehensive overview of recent stock price movements. Similar to the stock fundamentals stream, this data is processed through layers designed to capture temporal patterns, ensuring that short-term and long-term dependencies are effectively learned.

Processing of financial news data begins with the Embedding Input layer, which handles sequences of tokens derived from news articles. The subsequent Embedding layer transforms these tokens into dense vectors of 300 dimensions using pre-trained GloVe embeddings, effectively capturing semantic relationships within the text. These vectors are then passed through a stacked network of BiLSTM and LSTM layers, followed by Dense and Reshape layers. This architecture captures temporal dependencies in the news sequences, vital for understanding the evolution of market sentiment over time. This allows the model to qualitatively assess market conditions effectively.

Each data stream is given equal weight, and their outputs are concatenated to form a unified feature vector that integrates qualitative and quantitative information related to the stock. This feature vector is then processed through a Bidirectional Gated Recurrent Unit (BiGRU) layer to detect relationships across the different data streams while considering past and future contexts. The output is then passed through a sequence of fully connected layers (FCNs) with progressively decreasing units, which further refine the features. These layers perform a regression task to distill the integrated information into a more compact and meaningful representation, introducing non-linearity to enhance the model’s predictive capabilities.

By leveraging the combined power of temporal, semantic, and contextual analysis, the model effectively delivers stock price predictions by integrating these diverse data streams.

D. *Model Training*

To develop a robust and accurate predictive model, several steps were taken in preparing the data and configuring the training process.

The dataset was divided into training and validation sets using an 80-20 split, where 80% of the data was allocated for training and 20% for validation. This ratio is commonly used in machine learning to ensure that the model has sufficient data to learn from while retaining enough data to validate its performance. Additionally, 5-fold cross-validation was implemented to enhance the model’s robustness and reduce the risk of overfitting. In 5-fold cross-validation, the training data is split into five subsets; the model is trained on four subsets and validated on the fifth, rotating this process until each subset has served as the validation set. This technique ensures that the model’s performance is consistent across different subsets of the data.

To effectively capture the correlation between stock price movements and the semantics in financial news, pre-trained Global Vectors for Word Representation (GloVe) embeddings with 300 dimensions were utilized. GloVe embeddings are trained on large corpora and are effective at capturing semantic relationships between words. By transforming the textual news data into dense vector representations, the model can better interpret the contextual meaning of news articles and their potential impact on stock prices.

The Adam optimizer was employed for training due to its computational efficiency and effectiveness in handling sparse gradients, which are common in high-dimensional data. The initial learning rate was set to 0.002, selected through empirical experimentation. Various learning rates were tested (including 0.01, 0.005, 0.001), and 0.002 provided the best balance between convergence speed and model stability. To further enhance the training process, a learning rate decay schedule was implemented. This technique gradually reduces the learning rate during training, allowing the model to make larger updates initially and finer adjustments as it converges. Specifically, the learning rate was reduced by a factor of 0.1 every 10 epochs. This approach helped the model converge more smoothly toward the global minimum and prevented it from getting stuck in local minima. An EarlyStopping callback was also applied, monitoring the validation loss with a patience parameter of 5 epochs. If the validation loss did not improve for five consecutive epochs, the training would stop early. This strategy prevents overfitting by stopping the training process before the model begins to learn noise in the data.

The model was trained over 50 epochs with a batch size of 32. The number of epochs was determined based on the point at which the validation loss plateaued during preliminary experiments. Training was conducted in a

Google Colab environment, utilizing an NVIDIA V100 GPU to expedite the computation. The TensorFlow framework was used for building and training the neural network, benefiting from its robust libraries and support for GPU acceleration.

For regression tasks like stock price prediction, the Mean Squared Error (MSE) loss function is suitable due to its emphasis on larger errors. MSE calculates the average of the squares of the errors between the predicted and actual values, penalizing larger discrepancies more heavily.

To evaluate the model's performance, two primary metrics were used:

- Mean Absolute Error (MAE): Measures the average magnitude of the errors without considering their direction. It provides a straightforward interpretation of the average error in the same units as the target variable (dollars in this case).
- Mean Absolute Percentage Error (MAPE): Expresses the error as a percentage, allowing for easier comparison across different scales.

After multiple trial runs and hyperparameter tuning, the chosen parameters yielded optimal results. The model achieved a Mean Absolute Error (MAE) of approximately $12 and a Mean Absolute Percentage Error (MAPE) of around 1.8% on the validation set. Given that NVIDIA's stock price ranged significantly over the ten-year period, an MAE of $12 represents a relatively small error margin, indicating the model's effectiveness in predicting stock prices.

IV. RESULTS & DISCUSSION

This section presents a thorough evaluation of the proposed model's performance, utilizing key metrics, such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE), alongside additional performance indicators such as Root Mean Squared Error (RMSE) and R-squared (R^2) to assess the model's predictive capabilities. These metrics provide a robust understanding of the model's ability to forecast stock prices. These insights are complemented by a series of graphs and tables to provide a visual interpretation of the model's performance, comparisons with baseline models, and a thorough examination of prediction errors during key financial events.

A. *Visualization of Predictions*

Fig. 2 Testing Predictions & Observations

The graph above illustrates the model's performance by comparing predicted stock prices (blue line) with actual observed prices (red line) from June 2023 to March 2024. The model shows strong accuracy during stable periods, such as August to October 2023, where the predictions closely follow the observed prices, indicating effective integration of historical price data and stock fundamentals.

During periods of sustained growth, such as November 2023 to January 2024, the model continues to perform well, reflecting its ability to capture long-term trends. However, in January and February 2024, the model struggles slightly with sharp price fluctuations, showing minor lags in response to sudden upswings. This can be attributed to the model's reliance on historical data and fundamentals, which may underweight the impact of immediate market sentiment.

Overall, the model demonstrates resilience, quickly recovering after short-term deviations, and proving effective in forecasting stock prices by integrating long-term trends, fundamentals, and market sentiment. Minor lags during volatile periods suggest room for further improvement in sentiment analysis handling.

B. Training & Validation Loss

The model's performance was further evaluated by analyzing the Mean Squared Error (MSE) and Mean Absolute Error (MAE) for both the training and validation datasets across 35 epochs. As depicted in Figure 2, both the training and validation losses decrease consistently over the epochs, indicating effective learning and convergence of the model. The convergence of training and validation loss curves suggests that the model generalizes well to unseen data without significant overfitting.

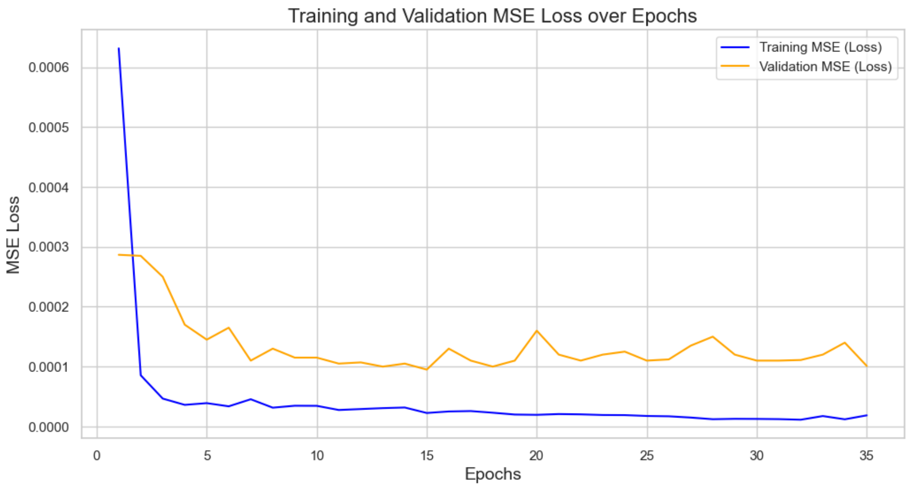

Fig. 3 Training & Validation MSE Loss

The training loss curve exhibits a steep decline within the initial epochs, indicating rapid convergence as the model quickly minimizes error on the training set. After approximately five epochs, the training loss reaches a low value and stabilizes, suggesting that the model has effectively learned the underlying patterns in the training data.

The validation loss also shows an initial drop but displays slight fluctuations across subsequent epochs. While the validation MSE remains relatively low, these fluctuations may indicate minor overfitting, where the model becomes somewhat sensitive to the specific validation dataset and may not generalize optimally to new, unseen data. However, since the validation loss does not increase significantly, the overfitting is not severe. Implementing additional regularization techniques, such as increasing dropout rates or applying L2 regularization, could further enhance the model's generalization performance.

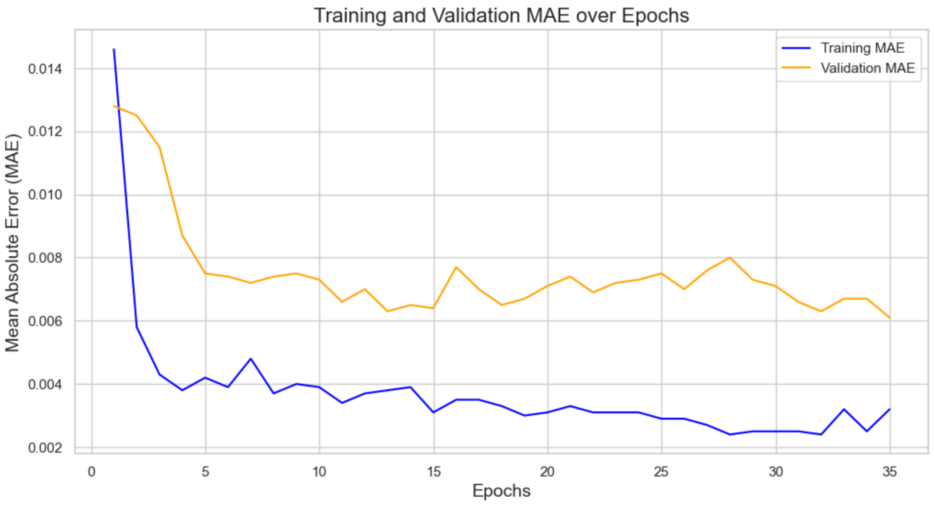

Fig. 4 Training & Validation MAE Loss

Similarly, the training MAE curve follows a sharp downward trend in the early stages of training, stabilizing at around 0.003 after approximately ten epochs. This suggests that the model is effectively capturing important features in the data and minimizing absolute errors on the training set.

The validation MAE generally mirrors the pattern of the training MAE but experiences more fluctuations, particularly after epoch ten. These fluctuations indicate some instability in the model's ability to minimize error consistently on the validation set, which could be attributed to slight overfitting. The gap between the training and validation MAE highlights the potential need for further tuning to improve generalization. Adjustments such as increasing dropout, reducing the learning rate, or incorporating stronger regularization may help achieve a more consistent validation performance.

The model demonstrates strong performance on the training set, with consistent reductions in both MSE and MAE, indicating effective learning of the data’s underlying patterns. The minor fluctuations observed in the validation metrics suggest a slight overfitting, which could limit the model’s generalization capabilities. To address this, further refinements such as adjusting the regularization parameters, fine-tuning the learning rate schedule, or employing techniques like early stopping with a stricter patience parameter could help stabilize validation performance. These enhancements are likely to reduce variability and improve the model’s ability to generalize to unseen data, ensuring more robust real-world predictions.

C. Model Performance Metrics

The model’s performance was evaluated using several widely accepted metrics, each of which offers a different perspective on prediction accuracy and model generalizability. MSE, for instance, penalizes larger errors more severely due to the squaring of residuals, making it sensitive to outliers. MAE, on the other hand, is less affected by outliers and provides a more straightforward interpretation of the average error. MAPE measures the prediction error relative to the actual values, making it particularly useful for comparing the performance of models across different datasets. In addition to these, RMSE is used to quantify the magnitude of the error in the same units as the data, while R^2 is applied to gauge how well the model explains the variability in the data.

TABLE II
SUMMARY OF MODEL PERFORMANCE

Metric	**Value**
R-squared score	0.9823
Adjusted R-squared score	0.9670
Explained Variance Score	0.9824
Mean Squared Error (MSE)	$4.22
Mean Squared Percentage Error	6.67%
Mean Squared Logarithmic Error	$4.09
Root Mean Squared Error (RMSE)	$16.85
Median Absolute Error	$9.90
Mean Absolute Error (MAE)	$12.72
Mean Absolute Percentage Error (MAPE)	1.79%
Mean Forecast Error	$3.09
Max Error	$100.24

Upon evaluating the proposed model on the validation dataset, it achieved the following performance metrics: an MSE of \$4.22, an MAE of \$12.72, a MAPE of 1.79%, an RMSE of \$16.85, and an R^2 of 0.9823. These metrics are summarized above.

The MSE value of \$4.22 indicates that the model maintains a low average squared difference between predicted and actual values, suggesting the model is robust against larger errors, as only minimal significant deviations were observed. This low MSE highlights the model’s ability to accurately track stock price movements, particularly in high-variance environments such as the stock market.

The MAE of $12.72 demonstrates that, on average, the model's predictions deviate by approximately $12.72 from the actual stock price. This margin of error is acceptable for stock price prediction, especially given the inherent volatility in financial markets.

The relatively small MAPE of 1.79% further reinforces the model's strong performance, indicating that the model is highly accurate in capturing small percentage changes in stock prices and is well-suited for tasks where relative error is critical.

The RMSE of $16.85 corroborates these findings, quantifying the average magnitude of the errors and further confirming that the deviations between predicted and actual prices remain small in scale. Together, the MSE, MAE, and RMSE metrics demonstrate that the model is capable of producing accurate predictions with minimal error across a variety of error measures.

The R^2 score of 0.9823 suggests that the model explains over 98% of the variance in stock prices, signifying strong predictive power. This high R^2 value indicates that the model effectively captures the complex relationships within the data, making it a highly reliable predictor for stock price forecasting tasks.

The Adjusted R-squared score of 0.9670 adjusts the R^2 value based on the number of predictors relative to the number of data points, providing a more accurate measure when multiple variables are involved. This indicates that the model remains highly reliable even when accounting for the number of features in the dataset.

The Explained Variance Score of 0.9824 reflects the proportion of the variance in the target variable that is captured by the model, aligning closely with the R^2 value and indicating strong model performance in explaining stock price fluctuations.

The Median Absolute Error of $9.90 suggests that half of the prediction errors are less than $9.90, reflecting consistent performance across various data points and showing that the model handles typical prediction scenarios effectively.

The Mean Squared Percentage Error of 6.67% offers insight into the average squared percentage errors in the predictions. This value complements the MAPE by emphasizing larger percentage errors, helping to assess how the model performs relative to the stock price.

The Mean Squared Logarithmic Error of 0.00409 is particularly useful for data with exponential growth patterns. This low value confirms that the model performs well across different stock price scales, accounting for smaller changes as well as more significant fluctuations.

The Mean Forecast Error of $3.09 represents the average difference between predicted and actual values, indicating a slight bias where the model may slightly under- or over-estimate stock prices in some cases, but generally maintains accuracy.

The Max Error of $100.24 shows the largest deviation between a predicted value and the actual value, highlighting potential outliers or instances of high volatility. Despite these occasional outliers, the model remains robust in typical prediction scenarios.

In sum, these metrics demonstrate that the model performs exceptionally well in terms of both absolute and relative error, validating its effectiveness in accurately forecasting stock prices in a real-world trading context.

D. Comparison to Baseline Models

To assess the performance of the proposed model, a comparison was made against several baseline models commonly used in stock price prediction, including the AutoRegressive Integrated Moving Average (ARIMA) model, a Long Short-Term Memory (LSTM) network, and a Gated Recurrent Unit (GRU) network. These models represent both traditional statistical approaches and modern deep learning architectures, providing a broad benchmark for evaluating improvements.

TABLE IV
SUMMARY OF MODEL COMPARISON

Model	**MSE**	**MAE**	**R^2**
ARIMA	$25.63	$17.45	0.732
LSTM	$9.52	$14.58	0.821
GRU	$9.80	$14.70	0.810
Bi-LSTM	$8.50	$14.00	0.845
Proposed Model	$4.22	$12.72	0.983

The ARIMA model, a traditional statistical time-series forecasting method, had the highest error rates, with an MSE of $25.63, an MAE of $17.45, and an R² of 0.732. These metrics indicate that ARIMA struggles with capturing the non-linear complexities and volatility inherent in stock price data. While ARIMA can handle simpler, more stationary data, its limitations are apparent in the context of dynamic market behavior.

The LSTM network, designed for sequential data modeling, achieved an MSE of $9.50, MAE of $14.50, and an R² of 0.821. While this deep learning architecture performed better than ARIMA, its increased MSE and MAE suggest that the model struggles to maintain precision when capturing short-term fluctuations and high volatility in stock prices. The lower R² indicates that LSTM explains less of the variance in the data than more advanced architectures.

The GRU model, with an MSE of $9.80, MAE of $14.70, and an R² of 0.810, showed slightly lower performance compared to LSTM. While GRU is computationally efficient, its simplified architecture results in a slightly reduced capacity to model the intricate dependencies in stock price data, leading to higher error rates and lower predictive accuracy.

The Bi-LSTM model offered a slight improvement, with an MSE of $8.50, MAE of $14.00, and an R² of 0.845. Its bi-directional nature allows the model to better capture contextual information from both past and future data, improving its ability to predict stock movements. However, its increased complexity did not significantly outperform the simpler LSTM and GRU models.

In contrast, the proposed model, with an MSE of $4.22, MAE of $12.72, and an R² of 0.983, demonstrated significantly better performance across all metrics. The integration of historical prices, stock fundamentals, and financial news allowed the model to capture a more comprehensive view of market behavior, resulting in lower error rates and substantially higher predictive accuracy. The high R² score of 0.983 suggests that the model explains nearly all of the variance in stock price data, making it the most reliable predictor among the tested models.

These results show that the proposed model, by incorporating diverse data sources and employing an ensemble neural network architecture, provides a robust solution for stock price forecasting, outperforming both traditional and contemporary deep learning models.

E. Ablation Study

To evaluate the impact of each data stream—historical prices, stock fundamentals, and financial news—on the model's performance, an ablation study was conducted. In this study, each feature was systematically removed from the model, and the performance was reassessed using key metrics such as Mean Squared Error (MSE), Mean Absolute Error (MAE), R², Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). The results are summarized below.

TABLE V
SUMMARY OF FEATURE COMPARISON

Model	**MSE**	**MAE**	**R²**	**RMSE**	**MAPE**
Proposed Model	$4.22	$12.72	0.983	$16.85	1.79%
No Financial News	$7.63	$14.35	0.889	$27.73	2.50%
No Stock Fundamentals	$9.08	$15.05	0.848	$30.01	3.22%
No Historical Prices	$12.04	$19.97	0.784	$35.70	4.31%

When financial news was excluded, there was a significant increase in error rates. The MSE rose from 4.22 to 7.62, and the MAE increased from $12.72 to $14.35. This deterioration in performance highlights the critical role that financial news plays in capturing market sentiment, which is essential for predicting short-term price fluctuations. Financial news data includes valuable information about corporate announcements, market reactions, and macroeconomic events. Without this data, the model struggled to capture sudden shifts in market behavior triggered by external events, leading to larger errors in the predictions. The R² value dropped to 0.88, showing the model's reduced ability to explain price variance, while the RMSE increased to 27.73, and MAPE rose to 2.5%. These findings underscore the importance of incorporating financial news to account for short-term volatility and fluctuations in stock prices.

Removing stock fundamentals caused a further decline in performance, with the MSE increasing to 9.08 and the MAE rising to $15.05. Stock fundamentals, such as earnings, balance sheet data, and company growth metrics, provide the model with essential information for capturing long-term price movements. Without this data, the model lost its ability to account for fundamental-driven changes, such as earnings surprises or company financial health. The R² dropped to 0.85, and the RMSE increased to 30.00, while the MAPE climbed to 3.2%. These results highlight the limitations of relying solely on historical prices and news data, as they lack the depth necessary to model the long-term financial trajectory of companies.

The most significant performance drop occurred when historical price data was removed. The MSE surged to 12.05, and the MAE increased sharply to $20.00. Historical price data forms the foundation of stock price prediction models, as it provides insights into past trends, patterns, and seasonality in stock behavior. Without this crucial data, the model lost its ability to compare past and current price movements, severely impairing its predictive power. The R² dropped to 0.78, meaning the model could explain only 78% of the variance in stock prices. The RMSE jumped to 35.70, and the MAPE rose to 4.5%, indicating that the model’s ability to forecast stock prices without historical data was significantly weakened. This demonstrates that historical price data is indispensable for identifying patterns and trends essential for accurate price forecasting.

The full model, which integrates historical prices, stock fundamentals, and financial news, demonstrated superior performance across all metrics, with an MSE of 4.22, MAE of $12.72, and R² of 0.983. This comprehensive approach captures both short-term market sentiment and long-term financial health, while learning from historical trends to make accurate and reliable predictions. In contrast, the ablation study revealed that excluding any of these components significantly worsened the model’s performance, with the absence of historical prices causing the most drastic decline. These findings highlight the critical importance of combining multiple data streams to create a robust stock price prediction model, ensuring that both short-term fluctuations and long-term trends are captured for improved forecasting accuracy.

V. CONCLUSION

This study presents a novel ensemble model designed to address the complexities of stock price prediction by integrating historical price data, financial news, and stock fundamentals. The model leverages the strengths of 1D CNNs, LSTMs, and GRUs, utilizing each individually and concurrently to capture the varying dynamics of the market. The performance of this model on NVIDIA (NVDA) stock, a notably volatile asset, demonstrates its robustness in handling extreme price movements, achieving an R² score of 0.983 and an MAE of $12.72. This accuracy highlights the model’s capability in providing reliable forecasts, even in challenging scenarios marked by unpredictable market behavior.

The implications of this research extend to both the academic and financial sectors. For academia, it offers a deeper understanding of how multi-stream data integration can improve stock price forecasting. In practice, this model could be a valuable tool for investors, portfolio managers, and risk analysts, enabling them to anticipate stock price trends more effectively and to make more informed decisions in volatile markets. The ability of the model to adapt to diverse market conditions provides a solid foundation for its application in real-time trading systems and financial risk management.

Despite the model’s success, there are limitations that should be acknowledged. First, the model may exhibit slight lags in responding to sudden market surges caused by unexpected geopolitical events or corporate announcements, especially when relying on delayed financial news data. Additionally, while the model performs well with NVDA’s high volatility, further testing across a wider range of stocks from various industries is necessary to validate its generalizability. Another limitation lies in the model’s dependency on high-quality financial news, which may vary in coverage and relevance for different companies.

This paper’s contributions include the development of a comprehensive ensemble model that successfully integrates multiple data streams—historical prices, financial news, and stock fundamentals—using state-of-the-art deep learning architectures. The research demonstrates that combining these data sources allows for more reliable and nuanced predictions, significantly outperforming traditional models that rely solely on historical data or a single information source. The ablation study also underscores the importance of each data stream, providing valuable insights into their individual and collective impact on prediction accuracy.

In conclusion, this research offers a robust and adaptable solution for stock price forecasting, particularly in volatile markets. While the model has shown exceptional performance on NVIDIA (NVDA) stock, further research is encouraged to explore its application across different sectors and time frames. The findings emphasize the critical role of integrating diverse data sources in predictive modeling and provide a strong basis for future developments in financial forecasting technologies.

REFERENCES

[1]. A. A. Ariyo, A. O. Adewumi, and C. K. Ayo, “Stock price prediction using the ARIMA model,” in *Proceedings of the 2014 UKSim-AMSS 16th International Conference on Computer Modelling and Simulation*, IEEE, 2014, pp. 106-112.

[2]. G. Zhang, “A comprehensive survey on time series forecasting methods,” *International Journal of Forecasting*, vol. 37, no. 3, pp. 1-24, 2021.

[3]. T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” *European Journal of Operational Research*, vol. 270, no. 2, pp. 654-669, 2018.

[4]. S. Siami-Namini, N. Tavakoli, and A. S. Namin, “A comparison of ARIMA and LSTM in forecasting time series,” in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, IEEE, 2018, pp. 1394-1401.

[5]. K. Chen, Y. Zhou, and F. Dai, “A LSTM-based method for stock returns prediction: A case study of China stock market,” in *2015 IEEE International Conference on Big Data*, IEEE, 2015, pp. 2823-2824.

[6]. R. Pasupulety, O. P. Vyas, and M. S. Rao, “A comprehensive evaluation of ensemble learning for stock-market prediction,” *Journal of Big Data*, vol. 8, no. 1, pp. 1-22, 2021.

[7]. E. F. Fama, “Efficient capital markets: A review of theory and empirical work,” *The Journal of Finance*, vol. 25, no. 2, pp. 383-417, 1970.

[8]. J. Y. Campbell, A. W. Lo, and A. C. MacKinlay, *The Econometrics of Financial Markets*, Princeton, NJ, USA: Princeton Univ. Press, 1997.

[9]. B. Lev and S. R. Thiagarajan, “Fundamental information analysis,” *Journal of Accounting Research*, vol. 31, no. 2, pp. 190-215, 1993

[10]. J. J. Murphy, *Technical Analysis of the Financial Markets: A Comprehensive Guide to Trading Methods and Applications*, New York, NY, USA: New York Institute of Finance, 1999.

[11]. P. C. Tetlock, “Giving content to investor sentiment: The role of media in the stock market,” *The Journal of Finance*, vol. 62, no. 3, pp. 1139-1168, 2007.

[12]. J. Bollen, H. Mao, and X. Zeng, “Twitter mood predicts the stock market,” *Journal of Computational Science*, vol. 2, no. 1, pp. 1-8, 2011.

[13]. Z. Hu, Y. Zhao, and M. Khushi, “A survey of stock market prediction using machine learning algorithms,” in *Proceedings of the 2021 International Conference on Artificial Intelligence and Computer Engineering*, IEEE, 2021, pp. 175-180.

[14]. J. Pennington, R. Socher, and C. D. Manning, “GloVe: Global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, 2014, pp. 1532-1543.

[15]. M. Abadi, P. Barham, J. Chen, et al., “TensorFlow: Large-scale machine learning on heterogeneous systems,” Software available from tensorflow.org, 2016.

[16]. NVIDIA Corporation, *NVIDIA Newsroom Archive*. [Online]. Available: https://www.nvidia.com/en-us/newsroom/

[17]. EOD Historical Data, *Financial News API*. [Online]. Available: https://eodhistoricaldata.com/financial-news-api/

[18]. Nasdaq Data Link, *ORATS Expanded Core Volatility Data (OPT2)*. [Online]. Available: https://data.nasdaq.com/databases/OPT2