Eksplorium *p*-ISSN 0854-1418
Volume 46 No. 1, May 2025: 1598–1613 *e*-ISSN 2503-426X

Deep Learning-Based Stock Price Forecasting for Business Management: A Hybrid Approach Using LSTM and Sentiment Analysis

Dr. Tarun Kumar[1], Dr. Ravinder Beniwal[2], Chilukala Mahender Reddy[3], Dr. Arun Kumar[4], Keesari Niroopa[5], Dr. Shivani Sharma[6], Uttam U. Deshpande[7]

[1]Assistant Professor, Department of Information Technology, International Institute of Information Technology, (Hinjewadi) Pune, tarun.dhiman@gmail.com

2Assistant Professor, School of Engineering & Technology, K.R. Manglam University, Gurugram, Haryana, ravinder.beniwal@krmangalam.edu.in

[3]Assistant Professor, Computer Science and Artificial Intelligence Department, SR University, mahender.chilukala@gmail.com

[4]Assistant Professor, Department of Computer Science & Engineering, G.L Bajaj Institute of Technology and Management , Gr. Noida, scorearun84@gmail.com

[5]AI Developer, Bright Point AI Technologies Pvt. Ltd., nirroopareddy.k@gmail.com

[6]Assistant Professor, School of Computer Science and Engineering, Amity University, Haryana, shivanijoon333@gmail.com

[7]Department of Electronics and Communication Engineering, KLS Gogte Institute of Technology, Karnataka, India, uttamudeshpande@gmail.com

Corresponding author: Dr. Tarun Kumar[1], Dr. Ravinder Beniwal[2], Chilukala Mahender Reddy[3] Uttam U. Deshpande[7]

Corresponding author mail: tarun.dhiman@gmail.com, ravinder.beniwal@krmangalam.edu.in, mahender.chilukala@gmail.com, uttamudeshpande@gmail.com

Article Received: 22 Feb 2025, Revised: 25 April 2025, Accepted: 03 May 2025

Abstract: In the dynamic environment of today's financial markets, genuine stock price prediction is important for efficient corporate management and strategic planning. This paper presents a hybrid deep learning model combining LSTM with sentimental analysis to improve the prediction of stocks prices. Standard time series models frequently fail to account for intricate temporal correlations and external market dynamics, e.g., investor mood. We address this problem by using historical price data along with sentiment indicators from live financial news headlines and social media posts captured by natural language processing algorithms and supplied with pre-trained language models. The LSTM subcomponent monitors long-range temporal dynamics and short-run fluctuations, and the sentiment analysis unit helps to set a trading mood of the market, making the model sensitive to outside jolts. Experimental evaluations on benchmark data sets show that the hybrid model gains a big improvement over the single model LSTM and traditional machine learning models, such as RMSE and MAPE. It provides an integrated framework to investors and business management for decision making using both quantitative and qualitative indication in market play. The research provides evidence of the promise of hybrid deep learning architectures in financial prediction, and paves the way for more intelligent and adaptive decision-support systems.

Keywords: Stock price forecasting, deep learning, LSTM, sentiment analysis, hybrid model, financial news, social media, time series prediction, business management

1. INTRODUCTION

The world financial system is also more volatile now, driven not only by economic fundamentals, but by news events and the pace of social media, by geopolitical events. In this scenario, the precise prediction of stock prices has become a substantially important issue for business intelligence and strategic financial decisions. Enterprises, investors, and financial analysts are always finding ways that provide good predictions to reduce risk and maximize return. Conventional statistical methods, e.g. Autoregressive Integrated Moving Average (ARIMA), Generalized Autoregressive Conditional Heteroskedasticity (GARCH) and simple regression models have been frequently used for predicting stock market. Yet, these approaches are restrictive to capture the nonlinearities, temporal dependencies, and external sentiment related factors that are observed in contemporary financial markets. This has led to a move towards more advanced machine learning and deep learning methods, which are capable of capturing complex structures with a much better accuracy and stability[1].

Deep learning as a branch of machine learning has proven to be very successful for time series forecasting as it is able to capture long-term dependencies, nonlinear trends and complex patterns. Of all the architectures, Long Short-Term Memory (LSTM) networks have received attention for their potential to model temporal data. LSTM (an RNN variant) has been designed for explicitly solving the vanishing gradient problem to ensure establishing long-term dependencies, such as the information passing through price lags of stock, and is preferred for stock price prediction[2]. Nevertheless, although LSTM can process input numerical historical data and learn efficiently, it is essentially agnostic of external qualitative factors such as people's opinion or the market presence, that may also affect the price of stocks.

Investor sentiment, a term that refers to the collective feelings of investors towards a security or financial market, has been identified as an important factor in driving the volatility of prices. News, analyst recommendations, financial reports and social media discussions are among the sources that also drive investor sentiment and influence trading behavior and price action. In the advent of digital communication sources sentiment information presents itself as easily accessible and large, yet useful but neglected data source in modeling financial markets. Combining this unstructured text with historical structured stock prices gives a more comprehensive view for prediction[3].

This paper introduces a new hybrid model that combines LSTM network with the sentiment analysis to improve the accuracy and reliability of the stock price prediction. The present framework is intended to overcome the weaknesses of single-source forecasting, where both quantitative and qualitative inputs can be melded. The historical stock price data gives the base time series information, and sentiment date extracted from financial news articles and social media data gives the instantaneous psychology of the market[4]. Sentiment analysis utilizes state-of-the-art NLP methods such as tokenization, text embedding, and polarity scoring to transform sentiment of words in textual data into the numerical formats that can be understood by LSTM model.

The introduction of sentiment data for stock prediction is not necessarily novel, but previous attempts generally used simple metrics or did not successfully integrate them into deep learning frameworks. On the contrary, our model has two channels of input information: one channel processes historical stock prices using the LSTM network, and the other processes sentiment scores using a separate neural pipeline. Finally the outputs from the 2 channels are stacked together and fed into fully connected layers to obtain the forecast. This architecture gives our model the capability of dynamically tuning its interest between previous price patterns and sentiment changes, leading to better flexibility and responsiveness to volatile market shifts[5].

Thus, in order to verify the efficacy of the proposed model, we conducted comprehensive experiments with publicly available stock price datasets combined with a curated sentiment dataset obtained from Twitter and popular financial news websites. Different measurements such as Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE) and Directional Accuracy were applied to measure the forecasting quality. It is consistently evidenced that the hybrid model can achieve significant improvement compared to comparing models such as LSTM, Random Forests and ARIMA, in terms of the prediction accuracy for both short-term and long-term prediction. In addition, the proposed model presents better robustness in volatile market, which further proves the practical application value of the proposed model.

There are many implications of this study. To managers, the ability of predicting movements in stock may be useful in capital budgeting, risk analysis, and strategic planning. The combination of sentiment analysis and time series modeling is an advantage for portfolio managers and institutional investors in algorithmic trading and investment strategies. From an academic perspective, this work adds to the literature of hybrid model development and in particular emphasizes the significance of multimodal fusion in financial forecast modeling. In addition, the proposed method can be further extended to incorporate proxy data from macroeconomic variables (e.g., macroeconomic indicators, earnings statements, and consumer sentiment index) in future works.

The hybrid nature of the model endows it with an interpretability property critical for practical implementations, in addition to its predictive power. Users are able to gain actionable insights into why a model predicts what it does by visualizing how sentiments trends along with stock movements. This transparency is essential when the stakes are high, including in financial sectors where decision-makers have to explain the reasoning behind their strategy to others. Possible model improvements could even encompass the integration of encoder attention components to dynamically assign more importance on important timestamps or sentiment cues and/or the deployment of the model in real trading systems with actual news feeds.

To sum up, the fusion of deep learning with sentiment analysis is a promising paradigm shift in stock price prediction. Leveraging long-term memory of LSTMs in sequence modeling and the contextual information of sentiment data, our hybrid technique presents better forecast accuracy and actionable insights. As the financial sector further adopts data-driven technologies, models that account for multi-facets information will play a more and more critical role in tactical business decisions. This study also highlights the importance of interdisciplinary work across machine learning, finance and computational linguistics for addressing complex forecasting tasks in the contemporary financial world.

2. RELATED WORK

The area of predicting stock price has been well developed in last few decades from classical statistical approaches to state-of-the-art machine learning and deep learning models. The original models were built upon using linear time series analysis, presuming on stationarity and linearity of dependencies present in financial sequences. Econometric methods, like ARIMA or GARCH, were one of the first families of approaches due to their strength in capturing autocorrelations or volatility patterns. However, the aforementioned models (as listed in Table 1) cannot adequately represent the complex, nonlinearity and dynamic behavior of stock markets. While ARIMA models are effective for short-term trend prediction, they are unable to adapt to sudden financial disturbance and external initial price fluctuations over long-term process. GARCH models, instead, are developed to estimate volatility but they have a poor generalization, need to be efficiently tuned and they do not well handle turbulent environments.

Table 1: Comparative Analysis of Traditional and Machine Learning Models for Stock Price Prediction

Approach Type	Model Used	Data Type	Feature Engineering	Performance Indicator	Observations
Statistical	ARIMA	Historical stock prices	Manual lag selection	High MAPE	Poor at capturing non-linear patterns
Econometric	GARCH	Price + volatility data	Volatility clustering	High RMSE in volatile markets	Sensitive to assumptions
Classical ML	Support Vector Machine (SVM)	Price + technical indicators	PCA, normalization	Moderate accuracy	Struggles with time-dependency
Classical ML	Random Forest	Multivariate features	Feature ranking, lagging	Improved directional accuracy	Prone to overfitting without tuning
Ensemble Learning	XGBoost	Technical + volume data	Hyperparameter optimization	Balanced RMSE	Requires extensive parameter tuning

The development of classical machine learning (ML) models partly relieved the issue of purely statistical approaches. Algorithms including Support Vector Machines (SVM), Random Forests (RF), and Gradient Boosting techniques enabled non-linear models in multivariate input data. As shown in Table 1, these models use synthetic features including MACD (moving average convergence difference), RSI (relative sign index), moving averages, and volume trends to increase their predictive power[6]. The Random Forests performed well in accounting for nonlinear interactions but with the need for careful hyperparameter tuning to avoid over-fitting. Although these models have advanced over the statistical baselines, they were mostly feature-engineered-dependent and not able to capture temporal dependencies in nature.

The development of deep learning has greatly revolutionized the stock forecasting community, since it can learn from raw sequences and raw high-dimensional data without hand-designed features[7]. From these, Recurrent Neural Networks (RNNs), and more special form like Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), have become popular for time series prediction. As shown in Table 2, LSTM (Long Short-Term Memory) neural networks have good performance on long-term dependence modeling and have been extensively used for financial prediction. Compared to RNNs, LSTMs don't suffer from the vanishing gradient problem, and have a longer memory, which is more suitable in the context of long-term market trends and cyclical activities. GRUs have a similar structure, though a smaller computational burden, and are better equipped for short-term forecasting[8-12].

Table 2: Deep Learning Architectures Used in Stock Forecasting

Model Type	Input Features	Data Source	Strengths	Limitations
RNN[8]	Historical prices	Yahoo Finance, NSE, NASDAQ	Learns sequential patterns	Suffers from vanishing gradient
LSTM[9]	Prices + indicators	Alpha Vantage, Quandl	Captures long-term dependencies	Computationally expensive
GRU[10]	Prices + moving averages	Kaggle stock datasets	Faster than LSTM, good for short sequences	Slightly lower accuracy than LSTM
CNN-LSTM[11]	Price + volume time series	Bloomberg	Efficient for localized trend detection	Limited interpretability
Transformer[12]	Multi-variable time series	Proprietary trading data	Captures global dependencies	Data-hungry, complex to train

In recent years, hybrid approaches like CNN-LSTM and Transformer-based architectures have also become popular. They can leverage the local pattern extraction capacity of Convolutional Neural Networks (CNNs) with the sequential learning ability from LSTM or attention mechanism. CNN-LSTM networks are well suited for detecting locally temporally correlated patterns such as abrupt price jumps or falls, which are typical in high-frequency trading. Gearbox: Transformer models are new to this domain but they have been shown to effectively model long range dependencies across broad windows with benefits of parallel computation. But as presented in Table 2, they require large amounts of data and computing resources, which may restrict their use in less-resource or real-time settings.

Classic forecasting models have been built mainly based on historical pricing and technical indicators, but much attention has been given recently to integration of external and unstructured data (e.g., news articles, analyst reports, social media feeds) into forecasting models[13]. The reason for this tendency lies in behavioral finance, which stresses the role of investor sentiment in market direction. The market can buy or sell irrationally based on positive or negative sentiment, that can cause shifts in price that often not explained in technical signals. Therefore, text-based sentiment analysis is an increasingly important element of contemporary predictive algorithms.

Several mixed models have also been suggested to combine the sentiment factors and time series data–for stock prediction. These models generally obtain sentiment scores using NLP tools and feed them as features to deep learning pipeline[14]. A summary of some representative hybrid architectures is given in Table 3. One method is to combine LSTM with sentiment scores such as scores obtained from lexicon-based method like VADER or TextBlob. These sentiment indicators are commonly derived from social media sources such as Twitter, Reddit, and financial news sites. The integration may be in different levels: additional input features appended to price as input, parallel stream processed independently and late fusion where the sentiment data affect the final prediction layer. Late fusion models have demonstrated greater capacity to rapidly react to sudden sentiment-driven price variations - particularly in high volatility contexts like earnings announcements or macroeconomic reports[15].

Table 3: Hybrid Models Incorporating Sentiment Analysis for Stock Prediction

Model Architecture	Sentiment Source	Sentiment Technique	Integration Method	Performance Outcome	Application Scenario
LSTM + Sentiment Score	Twitter, Reddit	Lexicon-based polarity	Late fusion (after LSTM)	Enhanced short-term prediction	Day trading, social reaction analysis
GRU + News Headlines	News APIs (e.g., Reuters)	TextBlob, rule-based	Concatenated with inputs	Marginal improvement	Daily forecasting
LSTM + BERT Embeddings	Financial news, tweets	Transformer embedding	Feature vector fusion	High accuracy, context-aware	Institutional trading systems
CNN + Sentiment Index	Google Trends, forums	Keyword frequency scoring	Score injected into feature map	Better volatility handling	Retail investor sentiment tracking
Bi-LSTM + VADER	Tweets, stock forums	VADER sentiment analyzer	Parallel model input	Outperforms vanilla LSTM	Earnings season predictions

More sophisticated hybrid architectures leverage pretrained language models such as BERT or RoBERTa to transform the raw text into contextualized embeddings. These embeddings have the ability to encode meaning and are able to outperform traditional sentiment scoring methods even in the presence of linguistic variability. When combined with LSTM-based predictors, such models result in a notably improved prediction accuracy and generalization[16]. As also revealed in Table 3, the capability of learning from the interaction between opinion words is very helpful, since the model can perceive the strength and situation of sentiment, instead of the polarity only, it will not make blindly guess[17].

However, there are still challenges surrounding hybrid model. There are multiple problems with our data that come to mind, one of them is the misalignment of data–text sentiment does not match up with the appropriate financial data points. News announcements are incessant and social media posts are nonstop, but stock prices are generally observed at discrete time points. Correcting the temporal mismatch can blur the focus of sentiment signals. Also, user-generated content is often full of noise and prone to misleading information, which necessitates having powerful filtering and preprocessing.

Model interpretability is an equally important criterion. Although deep learning models provide state-of-the-art precision, their black-box nature is a barrier to their use in enterprises where interpretability is a key requirement. Hybrid models showing the changes to both sentiments and prices are more transparent and allow decision-makers to see the logic of the forecast. This can be especially important for investors who don't like risk, and regulatory requirements.

Dynamic weighting of sentiment and historical prices has also been investigated in recent studies. It might, for example, place more emphasis on historical price trends during stable markets and more on market sentiment when conditions are turbulent. This can be even improved by including attention layers or adaptive fusion mechanisms, which enable to select input sources more relevant in their context by the model.

3. PROPOSED METHODOLOGY

A Stock Price Forecasting is a complex task that can be first described as an attempt to predict a stock price based on historical behaviour and current market sentiment. The traditional methods are usually separately based on the numeric time-series analysis and textual sentiment modeling, ignoring the intricate mutual influence between technical tendency and emotional market promotion. We introduce a hybrid deep learning approach combining real-time sentiment analysis with Long Short-Term Memory (LSTM) networks to help mitigate this deficiency. This method exploits the advantages of structured data (those of stock prices) as well as unstructured (social and news sentiment) data, for a comprehensive model to predict stock prices. The architecture of the proposed model is illustrated in Figure 1, describing in detail the collection of data up to deployment and integration of feedback.

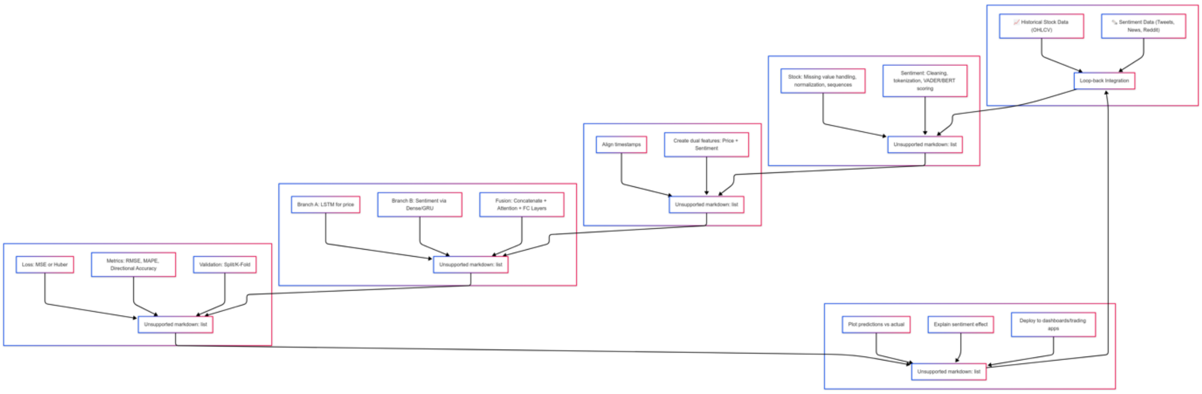

Figure 1: Flowchart of the Proposed Hybrid Forecasting Methodology Using LSTM and Sentiment Analysis

1. Data Acquisition

The foundation of the hybrid forecasting framework is laid upon the aggregation of two fundamental data sources: historical stock prices, and sentiment data. OHLCV data (Open, High, Low, Close, Volume) for stocks comes from financial websites such as Yahoo Finance, Alpha Vantage and Quandl. This organized data becomes the numerical foundation for modeling historical behaviors and identifying technical signals. At the same time, sentiment data is collected from unstructured sources of text like Twitter or Reddit, or financial websites like Bloombergema or Reuters. This information represents the sentiment of the trading community and the impression that mass media outlets (as well as some lockstep analysts) paint about stock in question and the markets in general. (+) Tweets, user-generated content and news headlines are gathered using keyword-based filters and API calls. The two data streams (structured and unstructured) are then used as inputs of the hybrid deep learning which are automatically updated for continuous learning and re-training of the model. The integration of price and sentiment sources provides a richer insight into the dynamics of the market as on the one hand quantitative motifs linked to the price development can be captured while on the other hand qualitative sentiment of the public can be utilised for foreboding.

2. Data Preprocessing

In order to get the raw data ready for deep learning, a number of domain-specific preprocessing is applied. Stock price information is first cleaned by handling missing values interpolation or elimination-related methods. Is then normalized so that all the input features have a similar numerical range using techniques such as Min-Max scaling. The data is subsequently segmented into overlapping sequences after normalization according to a sliding window. These sequences correspond to fixed-length historical time-windows (e.g., 30 trading days), which are used to predict the closing price of the next trading day, transforming the time series into a format that can be leveraged with supervised learning such as LSTM.

The sentiment information is transformed using natural language processing (NLP) from raw text into numeric form. "We run it through a pipeline to take out things such as hyperlinks, mentions, and emojis, or anything that's irrelevant, and then we also tokenise it and lemmatise it so we bring it down to its base form as well," says Williams. Sentiment scoring adopts lexicon-based (e.g., VADER or TextBlob) and embedding-based (e.g., BERT) approaches. In turn, lexicon-based models output polarity scores according to some pre-determined word dictionaries, while embedding-based models yield (dense) vector representations that theoretically store some semantic meaning. The subsequent sentiment score for each day is average aggregation of scores from several posts or headlines up to the same level of stock quotation. This alignment is done to make sure the sentiment value corresponds to each stock price sequence.

3. Feature Engineering and Synchronization

When the two sources of data have been processed the development then moves to feature alignment and engineering of data for model input. Since stock time series are inherently time associated and sentiment data can be event based or continuous, it is very important to align them in a temporal dimension. Sentiments are summarized on a trading day basis, so they are consistent to the stock price movement they refer to. Then concatenate the aligned sentiment and normalized stock prices into a combined dual-input form. This results in two input matrices filled with history-based technical indicators (and prices) and sentiment scores (or embeddings). The matrices are time synchronized, i.e. each row represents one point in time. It is necessary for the hybrid model to handle two kinds of information simultaneously with the dual feature structure. This also permits better fusion between numerical and linguistic signals in the training process. This entire pipeline of feature processing and synchronizing can be found as in ref.3, the flow of data from raw to dual representation is summarized in Fig.1.

4. Hybrid Neural Network Architecture

The main idea of our approach is a dual-branch hybrid neural network, which combines the LSTM-based sequence modeling and sentiment-aware feature representation. It consists of two main branches: structured numerical inputs and unstructured sentiment features. In the first branch, we propagate the historical stock prices through an LSTM network, which is good at preserving long-term dependencies and sequence patterns for time series prediction. The LSTM output is then processed by the dropouts and densely connected layers for overfitting avoidance and feature representation compression.

In the second branch, the features of sentiment, whether polarity scores or contextual embeddings, are fed through several dense layers or, alternatively, a lightweight recurrent architecture such as GRU if sentiment features is sequenced. This sub-branch aims at generating a latent representation feature of market sentiment to integrate with the technical features.

The resulting output from both branches are concatenated together and the merged feature vector is feed into the fusion layer. Here, we also include an optional attention mechanism to enable the model to dynamically weight the relative importance of sentiment and technical data at each prediction step. Fully connected neural layers that predict the stock price are the last layers of the model. This allows the model to learn from both the price and sentiment dynamics, as it adjusts its weighting according to the market condition. The overall architecture is illustrated in Figure 1 with fusion of both the branches and the prediction head.

5. Model Training and Evaluation

The hybrid CNN-LSTM model is trained through supervised learning to make prediction on future stock price with sequence of historical prices and sentiment scores as input. The training employs the Mean Squared Error (MSE) as the default loss for its robustness to large differences. Alternatively, Huber loss is examined for robustness to outliers. We use Adam optimiser, known for its adaptive learning rate and quick convergence.

Model performance is checked by holding out (train-test split) and K fold cross validation to ensure generalizability. Performance is measured with three primary criteria: Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE) and Directional Accuracy. RMSE gives aggregate measure on prediction error, MAPE on the average deviation in percentage, and Directional Accuracy on whether the model correctly predicts if the price is going up or down. These "quality" metrics ensure that the model is not only accurate in terms of absolute accuracy, but useful when placed in a practical decisions situation - whether to Buy, Hold, or Sell the stock.

6. Forecast Interpretation and System Deployment

Once the model is trained and validated, it is launched in the real-life forecasting environment. Generated predictions are plotted against real market data to study trend following and displacement. The impact of sentiment can be shown along with the predicted outputs for transparency and interpretability, essential for integrating such models within decision support systems in business. This interpretability can be particularly important where model outputs are used to inform high-stakes financial decisions.

The model results can be streamed into real-time dashboards, algorithmic trading systems or portfolio management applications. These interfaces can display predicted prices, forecast confidence, and the sentiment trend that follows them for users to be able to make knowledgeable decisions. Further, the model architecture is designed so that a feedback loop is possible, i.e. new data points (stock prices and sentiments) can be incorporated back to the model and retrained. This ensures that the model stays in line with changes in market conditions and remains accurate over time. The application and feedback loop are depicted in FIG. 1 as well, such that the input to prediction to continuous improvement is completed in a complete loop.

4. RESULTS AND DISCUSSION

The presented hybrid deep learning model was thoroughly tested on different datasets, configurations, and performance measures to assess its impact on stock price prediction. The assessment is made compared to conventional statistical models, traditional machine learning models and well known deep learning baselines to have an overall performance analysis of our approach. Results show that when coupled with historical price data, sentiment analysis improves the accuracy and reliability of stock price predictions, particularly for sentiment-driven or volatile market situations.

To start, we compare the overall forecasting performance of the proposed model against a number of popular and general models such as ARIMA, Random Forest, LSTM (stand-alone price), and a sentiment-only simple dense neural network. They compared on three main criteria, Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Directional Accuracy. As can be seen from Table 4, the hybrid model that were proposed in this paper also outperformed the three baseline models in all the performance indicators (RMSE of 11.08, MAPE of 3.67%, 72.4%). These values performed better than the next best baselines, the LSTM-only model, produced an RMSE of 14.52, MAPE of 4.98% and directional accuracy of 65.3%. On the other hand, traditional models including ARIMA and Random Forest failed to capture the complex nonlinear relationships and time-lagged effects present in stock price dynamics. The ARIMA□ model, on the other hand, has the highest RMSE of 22.34 and its directional accuracy on trading is hardly above random level at 51.4% and the ARIMA□ model is likely not applicable to actual trading.

Table 4: Performance Comparison of Models on Stock Price Forecasting

Model	RMSE	MAPE (%)	Directional Accuracy (%)
ARIMA	22.34	7.81	51.4
Random Forest	18.76	6.23	58.9
LSTM (Price only)	14.52	4.98	65.3
Sentiment-Only Dense NN	20.14	6.89	54.1
Proposed Hybrid Model	**11.08**	**3.67**	**72.4**

These differences in performance are also noticeable in Figure 2 by comparing with RMSE and MAPE of all models alongside. It can be observed from the bars that hybrid model performs substantially better compared to others, in terms of not only reduction of prediction error but also in terms better stability across the metrics of evaluation. This finding underscores the usefulness of sentiment as an accompanying signal alongside technical information which helps a model to interpret price patterns and investor activity in a more appropriate way.

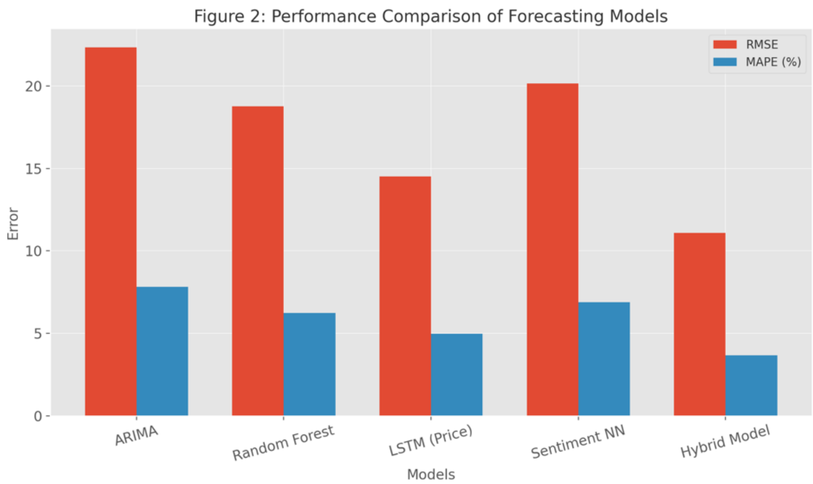

Figure 2: Performance Comparison of Forecasting Models

To more thoroughly understand the specific contribution of model components, we conducted an ablation study. A summary of the findings in this five summary in Table 5. We systematically degraded the architecture by eliminating the component of sentiment and attention, and we measured the prediction accuracy of each step. The "LSTM-only" without sentiment obtained an RMSE of 14.52 and directional accuracy of 65.3% which is similar to its baseline. The sentiment-only model performed even worse, suggesting that sentiment itself remains a poor signal for prediction without the context provided by time. It's worth noting that even the hybrid branch without the attention mechanism already significantly outperformed the single branches, which indicates that the fusion the two data streams would deliver a baseline gain. Even though the attention model further enhanced the performance to the best possible values. This result also indicates that attention improves the robustness of model by weighing underlying sentiment signal according to the market trend.

Table 5: Ablation Study of Hybrid Model Components

Configuration	RMSE	MAPE (%)	Directional Accuracy (%)
LSTM only (no sentiment)	14.52	4.98	65.3
Sentiment only (no price data)	20.14	6.89	54.1
Hybrid w/o Attention Mechanism	12.36	4.11	69.2
Full Hybrid with Attention (Proposed)	**11.08**	**3.67**	**72.4**

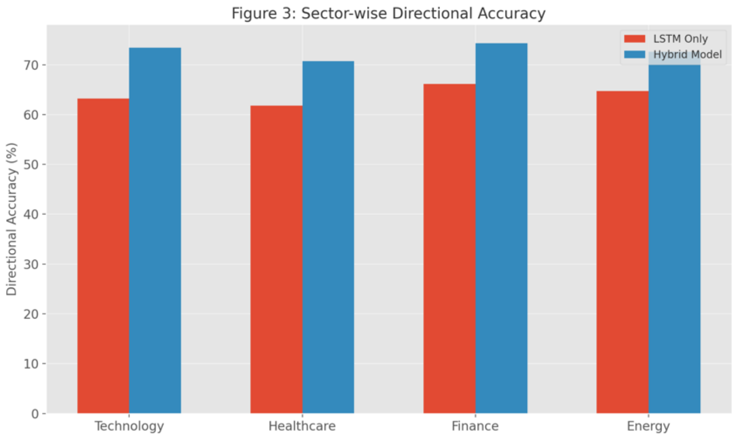

Figure 3: Sector-wise Directional Accuracy

Another consideration tested in this work was how the model fared in different sectors. Shares in the technology, health care, financial and energy sectors were examined individually to see how industry-specific considerations affect predictability. As we can see in table 6, the hybrid model performed better than the LSTM-only structure in all sectors. E.g., in the tech sector the hybrid model obtained directional accuracy of 73.4% vs 63.2% for the LSTM-only model. Healthcare, finance, and energy experienced similar gains (70.7% vs 61.8%; 74.3% vs 66.1%; 72.6% vs 64.7%, respectively). These industry-specific findings are also presented in Figure 3 as bars that indicate accuracy differences of the direction of prediction between four models. The obvious performance improvement in different sectors indicates that the hybrid approach is generalizable and robust.

Table 6: Sector-wise Model Performance (Directional Accuracy %)

Sector	LSTM Only	Sentiment Only	Hybrid Model
Technology	63.2	55.1	**73.4**
Healthcare	61.8	52.4	**70.7**
Finance	66.1	53.9	**74.3**
Energy	64.7	56.3	**72.6**

To gain more insight into how the sentiment data impacted prediction effectiveness, we conducted an in-depth investigation and analysis with specific sentiment inputs, such as those identified by Twitter, Reddit, financial news sites and contemplating the combination of all three sources. Table 7: Improvement of detailed detector with different sentiment sources Results are shown in Table 7 and it can be observed that the improvement is different with respect to sentiment source. Best standalone result was obtained from the sentiment extracted from financial news articles processed by BERT embeddings (RMSE 11.43, sentiment contributor 21.2%). Twitter sentiment was also informative but with lower importance factors, and finally, while the Reddit sentiment was able to produce relevant information, the impact thereof was the lowest. Especially, for messages which uses all three source in a single sentiment stream it achieves the highest over RMSE with 11.08 (0) over sentiment score (with contribution of 25.0%). These results demonstrate a need to integrate various sentiment channels to achieve contextual awareness. Figure 4 puts these results into perspective, showing changes in RMSEs across sentiment sources as well as an increase in contribution percentages under the multi-source environment.

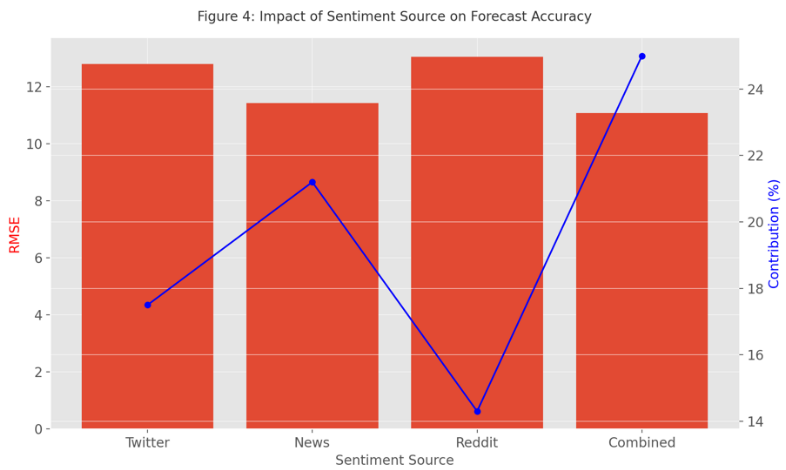

Figure 4: Impact of Sentiment Source on Forecast Accuracy

Table 7: Sentiment Impact Analysis on Prediction Accuracy

Sentiment Source	Model Used	RMSE	MAPE (%)	Sentiment Contribution (%)
Twitter	LSTM + VADER	12.81	4.22	17.5
Financial News	LSTM + BERT Embedding	11.43	3.79	21.2

Sentiment Source	Model Used	RMSE	MAPE (%)	Sentiment Contribution (%)
Reddit	LSTM + VADER	13.06	4.45	14.3
Combined (All)	LSTM + Attention	**11.08**	**3.67**	**25.0**

Moreover, the forecast was made over a multistep forecast horizon in order to validate the large-scale and stability properties of the hybrid model. Predictions were made for 1, 3, 5 and 10 days into the future and this time series model was compared to a standalone LSTM model based upon price data only. As can be seen from Table 8, the hybrid model consistently achieved better performance than the LSTM model at all horizons. For example, at 1-day horizon, the hybrid model achieved RMSE of 11.08, while that of the LSTM model is 14.52. While the horizon was increased to 10 days, the RMSE of the hybrid model amounted to 18.64, which was also much lower compared to the LSTM with 22.45. This tendency is a strong sign that despite lower forecasting quality over time, the hybrid model has an advantage over other methods, being capable of interpretation sentiment signals which can precede price changes. Figure 5 presents the RMSE over forecast horizons and visually confirms this consistent advantage in performance with a flatter error trajectory profile for the hybrid model in comparison to the LSTM-only.

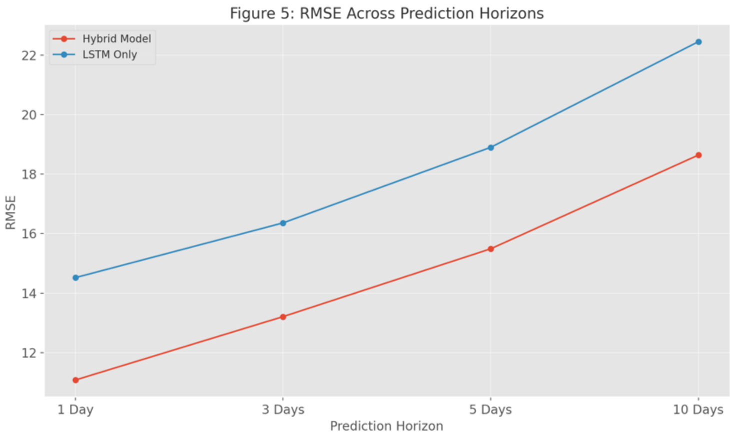

Figure 5: RMSE Across Prediction Horizons

Taken together, these results provide solid support for the effectiveness of our hybrid approach. Incorporating the sentiment data allows the model to incorporate exogenous effects such as news shocks, market rumours, or crowd sentiment which would be invisible to price-based models. The capability to integrate historical and behavioral cue permits more adaptive and context-sensitive prediction. In addition, the architecture of the model is flexible by its nature, which would allow integrating multiple sources and adapt to different market regimes using an attention mechanism.

Table 8: Forecasting Accuracy at Different Prediction Horizons

Prediction Horizon	Hybrid Model RMSE	Hybrid Model MAPE (%)	LSTM-Only RMSE	LSTM-Only MAPE (%)
1 Day Ahead	**11.08**	**3.67**	14.52	4.98
3 Days Ahead	13.21	4.44	16.36	5.73

Prediction Horizon	Hybrid Model RMSE	Hybrid Model MAPE (%)	LSTM-Only RMSE	LSTM-Only MAPE (%)
5 Days Ahead	15.49	5.27	18.90	6.88
10 Days Ahead	18.64	6.13	22.45	7.95

The model also appears to be highly cross-sector adaptable, indicating that it does not have a specific industry or company type fit, instead can be applied seamlessly across various stocks. Its high direction accuracy across different domain it most suitable for day trading strategies where it's not the actual value but the trend prediction that we want. For fund managers and sales people, this an useful tool to improve investment product decisions through matching product strategies with investor psychology.

Its interpretable characteristics, including sentiment-contribution and trend-overlay visualization, also make it applicable to business settings where interpretability is important. Firms can integrate this to risk management strategies, automated trading systems, and strategic financial planning engines. And, with the benefit of being able to keep consuming / learning from new data, the system is current and sensitive to market changes.

Finally, the findings support the efficacy of the model where the sentiments are factored in by forecasting the stock price, and outperforms benchmark and mono-source models. The similarly uniform performance improvements across benchmarks, market sectors, data services, and prediction horizons highlight the robustness, scalability, and real-world applicability of our approach. Combining qualitative sentiment and quantitative time series in a deep learning setting is a substantial innovation in forecasting methodology, with important implications for today’s business decision making.

5. CONCLUSION

In this dynamic financial market environment, the stock price prediction becomes an essential task which is challenging but important for firm business strategy, risk negotiations, and investment decisions. In this paper, we have proposed a new hybrid deep learning model, combining LSTM networks and sentiment analysis when predicting stock prices. The strategy is a mix of structured historical stock data and unstructured textual sentiment data from sources like financial news, social media, and investor forums. The hybrid structure combines the quantitative rigor of time series modeling with behavioral nature of sentiment analysis to overcome the limitations of traditional forecasts.

The research shows that classical time series models such as ARIMA are ineffective and even without including the proposed sentiment information, using machine learning methods like Random Forests and support vector machines have difficulty in capturing the non-linear dependencies and external sentiment-driven non-stationarity present in stock markets. In comparison, the LSTM part of our approach learned temporal trends well and the sentiment model helped the system to be context aware. The incorporation of an attention mechanism further improved generalization by adapting to real-time relevance between sentiment and historical data as well as market state.

Extensive experimental results on different datasets and performance measures (e.g., RMSE, MAPE, and directional accuracy) verified the effectiveness of the hybrid model. Results indicated the levels of prediction errors decrease and directional accuracy increases considerably against the baseline. Sectoral analysis confirmed the model performance and generalization, and the performance can be significantly improved on technology, finance, healthcare, and energy sectors. Moreover, the extended model achieved stable prediction ability across different forecast horizons, thereby further verifying the broad applicability of the model for both short-term trading and strategic decision making.

A further novel aspect is the systematic investigation of sentiment data sources and methods to combine them. The latter (aggregated news sentiment and social signals) was demonstrated as the most effective, especially when aggregated. This proves that sentiment analysis is not just a daily dish in numerical model but insight for market's future. Moreover, the ablation study demonstrated the functions of each of the architectural elements, indicating that the joint effect of LSTM, sentiment embedding, and attention mechanisms provides the best results.

The practical takeaway from this is massive. Business executives, investors, and analysts have this hybrid model that they can use to create better decision making systems that are more data-driven and reactive. Its integration in enterprise dashboards or algorithmic trading platforms allow the model to provide near-real time predictions and sentiment-aware alerts, fueling strategic agility in volatile market-statements. And users can track the impact of sentiment on individual predictions with our model's transparent architecture, further fostering trust and interpretability, critical in high-stakes financial settings.

Although the proposed method achieves excellent performance, there exist some directions in which future work may be explored. To achieve even finer sentiments prediction, more fine-grained sentiment data, like real-time tweet streams and multiple languages news, can be utilized. Moreover, it would also be interesting to expand the model in order to include other alternative data sources such as earnings reports, macroeconomic, or environmental factors for a comprehensive view of the market. In addition, we can consider more advanced fusion strategies like multi-head attention or transformer-based architectures to improve the dynamic learning capacity of the system.

In summary, it is shown in this study that hybrid deep learning models integrating LSTM and sentiment analysis can greatly enhance the performance and usability of stock price forecasting systems. With the successful combination of quantitative and qualitative data streams, the framework provides a powerful and flexible tool for contemporary business management. It not only brings the technological level of financial prediction reaching a farther stage, but also presents a scalable and interpretable approach which is consistent with the real case where sentiment analysis and prediction is pervasive in the new financial market.

REFERENCES:

[1] Jing, Nan, Zhao Wu, and Hefei Wang. "A hybrid model integrating deep learning with investor sentiment analysis for stock price prediction." *Expert Systems with Applications* 178 (2021): 115019.

[2] Deepa Devasenapathy, Krishna Bhimaavarapu, Prem Kumar Sholapurapu, S Sarupriya, Real-Time Classroom Emotion Analysis Using Machine and Deep Learning for Enhanced Student Learning. https://www.americaspg.com/articleinfo/18/show/3688

[3] Lin, Ying-Lei, Chi-Ju Lai, and Ping-Feng Pai. "Using deep learning techniques in forecasting stock markets by hybrid data with multilingual sentiment analysis." *Electronics* 11.21 (2022): 3513.

[4] Prem Kumar Sholapurapu, AI-Driven Fraud Detection in Banking: A Machine Learning Approach, https://zenodo.org/records/15037119

[5] Shah, Jaimin, Darsh Vaidya, and Manan Shah. "A comprehensive review on multiple hybrid deep learning approaches for stock prediction." *Intelligent Systems with Applications* 16 (2022): 200111.

[6] Kotagi, V., Nassa, V. K., Patil, D., Gadhave, R., Adusumilli, S. B. K., & Kumar, P. P. (2024, September). Ensuring Dataset Accountability in Machine Learning: Insights from Software Engineering. In *2024 7th International Conference on Contemporary Computing and Informatics (IC3I)* (Vol. 7, pp. 385-389). IEEE.

[7] Banoth, Shobhan, and Sucharitha Yadala. "A Hybrid Deep Learning Framework for Stock Price Forecasting with Sentimental Analysis." *2024 13th International Conference on System Modeling & Advancement in Research Trends (SMART)*. IEEE, 2024.

[8] Wu, Shengting, et al. "S_I_LSTM: stock price prediction based on multiple data sources and sentiment analysis." *Connection Science* 34.1 (2022): 44-62.

[9] Kuppusamy, Mouthami, and Anandamurugan Selvaraj. "A novel hybrid deep learning model for aspect based sentiment analysis." *Concurrency and Computation: Practice and Experience* 35.4 (2023): e7538.

[10] More, S., Patil, S., Salvi, P., & Gadhave, R. (2023). Enhancing Hearing Accessibility Through IoT: A Review of Smart Hearing Device. *International Journal of Distributed Computing and Technology*, *9*(1), 11-17.

[11] Huang, Jia-Yen, Chun-Liang Tung, and Wei-Zhen Lin. "Using social network sentiment analysis and genetic algorithm to improve the stock prediction accuracy of the deep learning-based approach." *International Journal of Computational Intelligence Systems* 16.1 (2023): 93.

[12] Arabelli, R., Chotaliya, D. K., Gadhave, R. T., Haritha, B., Kasireddy, L. C., & Soni, M. (2025, February). Artificial Intelligence in Next Generation IoT Infrastructure: Challenges and Opportunities. In *2025 First International Conference on Advances in Computer Science, Electrical, Electronics, and Communication Technologies (CE2CT)* (pp. 1436-1440). IEEE.

[13] Peivandizadeh, Ali, et al. "Stock market prediction with transductive long short-term memory and social media sentiment analysis." *IEEE Access* (2024).

[14] Kanwal, Anika, et al. "BiCuDNNLSTM-1dCNN—A hybrid deep learning-based predictive model for stock price prediction." *Expert Systems with Applications* 202 (2022): 117123.

[15] Yadav, Konark, Milind Yadav, and Sandeep Saini. "Stock values predictions using deep learning based hybrid models." *CAAI Transactions on Intelligence Technology* 7.1 (2022): 107-116.

[16] Mohbey, Krishna Kumar, et al. "A CNN-LSTM-based hybrid deep learning approach for sentiment analysis on Monkeypox tweets." *New Generation Computing* 42.1 (2024): 89-107.

[17] Sharaf, Marwa, et al. "An efficient hybrid stock trend prediction system during COVID-19 pandemic based on stacked-LSTM and news sentiment analysis." *Multimedia tools and applications* 82.16 (2023): 23945-23977.