

Stock market analysis and prediction

C.Sennilavazhagan

Sri Manakula Vinayagar Engineering College

Abstract:

Stock market prediction remains a challenging task due to the highly volatile, non-linear, and dynamic behavior of financial markets, where price movements are influenced by both historical trends and external factors such as news and investor sentiment. This paper presents a hybrid data-driven framework that combines Bidirectional Gated Recurrent Units (BiGRU) and Bidirectional Long Short-Term Memory (BiLSTM) networks for accurate stock price prediction, along with Natural Language Processing (NLP)-based sentiment analysis of financial news. The BiLSTM component is designed to capture long-term dependencies in time-series data, while the BiGRU enhances computational efficiency and learns short-term patterns effectively. The bidirectional structure enables the model to leverage contextual information from both past and future sequences, improving predictive performance. Simultaneously, financial news data is collected and processed using NLP techniques, including text preprocessing and sentiment extraction, to quantify market sentiment. These sentiment scores are integrated with historical stock prices and technical indicators to provide enriched input to the hybrid model. Experimental results indicate that the proposed approach significantly improves prediction accuracy and reduces forecasting errors compared to traditional models. The integration of sentiment-aware features enhances the model's robustness, making it suitable for real-time decision support in dynamic financial environments.

Keywords: *Stock Market Prediction, BiGRU, BiLSTM, Hybrid Deep Learning, Time Series Forecasting, Natural Language Processing, Sentiment Analysis, Financial News.*

I. Introduction:

The stock market broadly refers to the collection of exchanges and other venues where the buying, selling, and issuance of shares of publicly held companies take place. While both the terms “stock market” and “stock exchange” are often used interchangeably, the latter term generally comprises a subset of the former. If one trades in the stock market, it means that they buy or sell shares on one or more of

the stock exchange that are part of the overall stock market. A given country or region may have one or more exchanges comprising their stock market. The leading U.S. stock exchanges include the New York Stock Exchange (NYSE) and the NASDAQ. These leading national exchanges, along with several other exchanges operating in the country, form the stock market of the United States. Stock markets are venues where buyers and sellers meet to exchange equity shares of public corporations. Stock markets are vital components of a free-market economy because they enable democratized access to trading and exchange of capital for investors of all kinds. They perform several functions in markets, including efficient price discovery and efficient dealing. In the United States, the stock market is regulated by the Securities and Exchange Commission and local regulatory bodies. The process of stock market is shown in figure (1.1).



FIGURE 1.1 TRADITIONAL METHODS OF STOCK MARKETING

a. Hybrid model:

The combination of Bidirectional Gated Recurrent Units (BiGRU) and Bidirectional Long Short-Term Memory (BiLSTM) networks forms a powerful hybrid architecture for sequential data modeling. BiLSTM is effective in capturing long-term dependencies by maintaining memory over extended time steps, making it suitable for learning complex temporal patterns in stock price data. BiGRU, on the other hand, is computationally efficient and capable of quickly adapting to short-term fluctuations due to its simpler gating mechanism. By integrating both models, the system leverages

the strengths of BiLSTM in preserving historical context and BiGRU in capturing recent trends. The bidirectional nature of both networks allows learning from past and future sequences simultaneously, resulting in improved feature representation, enhanced prediction accuracy, and better generalization in dynamic environments.

b. Natural Language Processing (NLP)

Natural Language Processing (NLP) is a branch of artificial intelligence that enables machines to understand, interpret, and analyze human language in a meaningful way. It combines computational linguistics with machine learning techniques to process textual data efficiently. In the context of financial analysis, NLP is used to extract insights from news articles, reports, and social media content. The process typically involves steps such as text preprocessing, including tokenization, stop-word removal, and normalization, followed by feature extraction and sentiment analysis. NLP techniques can identify the polarity of text as positive, negative, or neutral, reflecting investor sentiment and market perception. By converting unstructured textual information into structured data, NLP helps improve decision-making and enhances the accuracy of predictive models in dynamic environments.

II. Literature Overview

[1] Stock market prediction with optimized PLSTM-AL in smart urban cities [1] Accurate stock market forecasting plays a vital role in guiding investors and market participants, as even slight improvements in predictive accuracy can lead to substantial financial gains. However, stock price forecasting remains a challenging task due to the inherent noise, complexity, and volatility in financial time series data. To address these challenges, this study introduces PLSTM-AL, a novel deep learning-based stock market forecasting model that combines Pareto-like sequential sampling optimization, Long Short-Term Memory (LSTM) networks, and an Attention Layer. The Pareto-like optimization component enhances data sampling by focusing on informative sequences, thereby improving model efficiency and reducing overfitting. The LSTM component, known for its ability to capture long-term temporal dependencies in sequential data, effectively learns intricate patterns within stock price movements. Furthermore, the integration of the Attention Layer enables the model to dynamically prioritize important time steps, allowing it to focus on the most relevant information that influences market direction. The PLSTM-AL model is evaluated using daily data from five major global stock indices covering the period 2003–2024. Its

performance is benchmarked against other advanced deep learning models such as Convolutional Neural Network (CNN) and Gated Recurrent Unit (GRU) networks. Experimental results show that PLSTM-AL consistently achieves higher accuracy and improved evaluation metrics across most datasets, confirming its superior robustness, adaptability, and predictive capability. Overall, the proposed PLSTM-AL framework demonstrates significant potential in enhancing stock market forecasting accuracy, offering investors a more reliable tool for decision-making in dynamic and uncertain financial environments. **[2] Combining market-guided patterns and mamba for stock price prediction** [2] Stock price prediction has long been recognized as a highly complex and challenging task due to the financial market's inherent volatility, nonlinearity, and dynamic behavior. With the rapid advancement of deep learning technologies, researchers have increasingly focused on modeling temporal dependencies to improve the accuracy of stock forecasting. Traditional methods typically utilize a shared neural network architecture to extract temporal patterns from individual stock time series and then merge these learned representations to infer relationships among different stocks. However, such approaches often fail to effectively capture fine-grained dependencies and complex inter-stock correlations, leading to suboptimal predictive performance. To address these limitations, a novel architecture called Market-Embedding with Mamba (MEM) is proposed. MEM introduces a multi-level feature aggregation framework consisting of coarse-grained, fine-grained, and temporal feature aggregation modules. The coarse-grained component captures broad market trends, the fine-grained component focuses on detailed stock-specific fluctuations, and the temporal aggregation module models dynamic dependencies over time. By integrating these three hierarchical levels of representation, MEM is capable of extracting more discriminative and informative features for precise stock price forecasting. Extensive experiments conducted on two widely used benchmark datasets—CSI300 and CSI800—demonstrate that MEM significantly outperforms existing state-of-the-art approaches. The superior performance of MEM highlights its robustness, adaptability, and ability to model intricate market dynamics effectively. Overall, the proposed MEM architecture represents a substantial advancement in the application of deep learning to financial forecasting, offering more reliable and insightful predictions for investors and analysts. **[3] Integrating PCA with deep learning models for stock market Forecasting: An analysis of Turkish stocks markets** [3] Financial data such as stock prices represent rich time-series information that provides critical insights for investors and financial analysts. Analyzing these data patterns is

essential for understanding market behavior and predicting future price movements. However, stock market prediction remains an inherently complex challenge due to the presence of intense noise, nonlinear patterns, and high volatility, which obscure clear trends and make accurate forecasting difficult. Despite these challenges, this complexity also creates opportunities for investors to uncover valuable market signals. One of the widely used approaches for stock prediction is technical analysis, which relies on multiple indicators that transform historical price data into actionable signals such as “buy,” “sell,” or “hold.” In this study, ten of the most commonly used technical indicators were analyzed using Principal Component Analysis (PCA) to reduce dimensionality and identify the most informative features. The main objective of the study is to explore the integration of PCA with deep learning architectures for predicting trends in the Turkish stock market, assessing how this combination impacts prediction performance. The PCA-selected indicators were used as inputs to develop four deep learning-based models employing architectures such as Long Short-Term Memory (LSTM) and Convolutional Neural Network (CNN). Model performance was evaluated using standard metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and R^2 (Coefficient of Determination). Experimental results demonstrate that integrating PCA with deep learning significantly enhances prediction accuracy. Among the proposed models, the PCA-LSTM-CNN hybrid model achieved the best results, indicating its strong capability in capturing both temporal and spatial dependencies in financial data, thereby improving the effectiveness of stock market forecasting. [4] **Sentiment analysis for stock market research: A bibliometric study [4]** Sentiment analysis has become an increasingly popular and influential approach in stock market research, as it provides valuable insights into investors' emotions, opinions, and market behavior. To gain a comprehensive understanding of research trends in this field, a bibliometric analysis was conducted on 223 scholarly articles related to sentiment analysis in stock market studies, published between 2010 and 2022, using data collected from the Web of Science database. This analysis examined various aspects, including active affiliations, contributing countries and regions, leading publication sources, and major subject areas. It also identified the most highly cited research works, visualized scientific collaborations among authors, institutions, and regions, and revealed prominent research themes. The results indicate that computer science journals play a central role in disseminating studies that integrate sentiment analysis with stock market forecasting. Moreover, research in this domain has attracted contributions from a wide range of geographical regions,

demonstrating global interest and collaboration. However, the study found that intra-regional collaborations—within the same geographical area—are stronger and more frequent than inter-regional ones. Using keyword mapping, key thematic areas were identified, including deep learning for stock market prediction, financial news sentiment-based stock trend forecasting, the effects of investor sentiment on market movements, and microblog sentiment classification for market analysis. These themes reflect the growing emphasis on combining sentiment analytics with machine learning to improve prediction accuracy. Overall, the findings provide a clear depiction of the research landscape, helping scholars and practitioners understand current progress, identify emerging trends, and strategically plan future studies on the application of sentiment analysis in financial markets. [5] **A Novel Stock Market Prediction Approach Based on Stacking Ensemble Machine Learning for Effective Trend Analysis [5]** The rapid economic growth in recent years has led to a substantial influx of capital into the stock market, making accurate prediction of market behavior an increasingly important and profitable endeavor. However, due to the dynamic and unpredictable nature of stock prices, forecasting remains a major challenge for investors. To address this issue, the Stacked Ensemble Model for Trend Analysis (SEMP-TA) is proposed as a novel and effective approach for predicting stock market trends. The SEMP-TA model operates in two primary stages: data preprocessing and trend prediction. During the preprocessing phase, the model performs data normalization, identifies and removes missing values, and employs cluster-based feature selection to rapidly determine the most influential financial indicators. This step ensures that only the most relevant and high-quality data are used for analysis. In the second stage, the cleaned and optimized data are fed into a stacked ensemble prediction framework that integrates three powerful machine learning algorithms: XGBoost, LightGBM, and Support Vector Regression (SVR). This combination leverages the strengths of each algorithm—XGBoost's gradient boosting efficiency, LightGBM's speed and scalability, and SVR's capability in handling nonlinear relationships—to enhance prediction accuracy and robustness. Experimental evaluations using the NIFTY-50 dataset demonstrate the superior performance of SEMP-TA, achieving a Mean Absolute Error (MAE) of 1.60, a Mean Squared Error (MSE) of 7.896, and an R^2 score of 0.8997. These results confirm that SEMP-TA outperforms existing state-of-the-art approaches in both accuracy and reliability, offering a highly effective solution for trend prediction and informed investment decision-making in volatile financial markets.

III. Methodologies and Approaches

a. Data Collection from Yahoo Finance

The data collection process begins with extracting historical stock market data from Yahoo Finance, a widely used and reliable financial data source. Yahoo Finance provides comprehensive market information, including Open, High, Low, Close, and Volume (OHLCV) values across various time intervals such as daily, weekly, or intraday data. This historical data serves as the foundation for understanding past market behavior and identifying price movement patterns. The data is collected programmatically using APIs or Python libraries, ensuring automation and consistency. By selecting appropriate time frames and stock symbols, the system gathers sufficient data to capture long-term trends as well as short-term volatility. Yahoo Finance is particularly suitable for this research due to its accessibility, accuracy, and wide coverage of global stock markets. The collected data is stored in structured formats, enabling efficient preprocessing and feature extraction in later stages. Reliable historical data is essential for training deep Data collection is the foundational step in any stock market prediction system, as the quality and reliability of input data directly influence model performance. In this proposed system, historical stock market data is obtained from Yahoo Finance, a widely used financial data platform that provides free access to comprehensive market information. Using APIs or Python libraries such as `yfinance`, essential stock attributes including opening price, closing price, high, low, adjusted close, and trading volume are retrieved over a specified time period. This data is collected for individual stocks or multiple companies, depending on the scope of analysis. Yahoo Finance also provides additional metadata such as dividends, stock splits, and company-related financial indicators, which can further enrich the dataset. The collected data is typically stored in structured formats like CSV or DataFrames for ease of processing. Ensuring data consistency, completeness, and correctness is crucial at this stage, as missing or noisy data can negatively impact predictions. Time-indexed data is especially important for time-series forecasting, as it preserves the sequential nature of stock price movements.

b. Pre-processing of Stock Market Data

Pre-processing is a critical step that prepares raw stock market data for effective use in machine learning models. The collected data often contains missing values, noise, and inconsistencies that must be addressed before training. Initially, missing values are handled using techniques such as interpolation or forward/backward filling to maintain continuity in time-series data. Outliers, which

may occur due to sudden spikes or data errors, are detected and treated to prevent distortion of model learning. The data is then normalized or scaled using methods like Min-Max scaling or Z-score normalization to ensure that all features contribute equally to the model, avoiding bias toward larger numerical values. Time-series data is further structured into sequences or sliding windows, where past observations are used to predict future values. Additionally, date-time features may be decomposed into components such as day, month, or trends to enhance learning. Proper data formatting is essential for feeding into deep learning models like BiGRU and BiLSTM, which require sequential input. Pre-processing not only improves model accuracy but also accelerates convergence during training. By transforming raw financial data into a clean and structured format, this step ensures that the predictive model can effectively capture meaningful patterns and relationships.

c. Feature Extraction

Feature extraction involves transforming raw stock data into informative attributes that enhance the predictive capability of the model. Instead of relying solely on basic price values, additional technical indicators are derived to capture market trends, momentum, and volatility. Commonly used indicators include Moving Averages (MA), Exponential Moving Averages (EMA), Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger Bands. These indicators help represent complex market behaviors in a structured form, making it easier for the model to learn patterns. For example, moving averages smooth out price fluctuations to highlight trends, while RSI indicates whether a stock is overbought or oversold. Volume-based features are also considered to understand trading activity and market strength. Feature extraction may also include lag features, where previous time steps are used as inputs for prediction. Dimensionality reduction techniques can be applied to remove redundant or irrelevant features, improving model efficiency. The extracted features are then combined into a feature vector that represents each time step in the dataset. This enriched representation allows the hybrid BiGRU-BiLSTM model to capture both temporal dependencies and market dynamics more effectively. Ultimately, feature extraction plays a vital role in improving prediction accuracy and model robustness.

d. Model Creation Using Temporal Convolutional Network (TCN)

The model creation phase involves designing and training a hybrid deep learning architecture that combines Bidirectional Long Short-Term Memory (BiLSTM) and Bidirectional Gated Recurrent Unit (BiGRU) networks. BiLSTM is used to capture long-term dependencies in stock price sequences by maintaining memory over extended time periods, which is essential for understanding historical trends. BiGRU, being a simplified version of LSTM, complements this by efficiently capturing short-term dependencies and reducing computational complexity. The bidirectional nature of both models enables the system to learn from both past and future contexts within the sequence, enhancing feature representation. In the hybrid architecture, BiLSTM layers are typically placed first to extract deep temporal features, followed by BiGRU layers to refine and optimize these features. The output is then passed through dense layers for final prediction of stock prices. The model is trained using historical data with appropriate loss functions such as Mean Squared Error (MSE) and optimization algorithms like Adam. Hyperparameters such as learning rate, batch size, and number of epochs are tuned to achieve optimal performance. This hybrid approach improves prediction accuracy, reduces overfitting, and ensures better generalization in dynamic financial environments.

e. News Collection Using Newspaper3K

News collection is an essential component for incorporating external factors that influence stock market behavior. In this system, financial news articles are collected using the Newspaper3K library, a Python-based tool designed for web scraping and article extraction. Newspaper3K automatically retrieves news content from various online sources, including financial websites, blogs, and news portals. It extracts important elements such as article title, publication date, author, and full text while removing unnecessary HTML tags and advertisements. This automation enables real-time collection of large volumes of news data related to specific companies or market sectors. The collected articles are filtered based on keywords such as company names, stock symbols, or financial terms to ensure relevance. Language detection and translation features can also be applied if needed. The extracted text is stored in a structured format for further processing. By continuously gathering up-to-date financial news, the system ensures that sentiment analysis reflects current market conditions. This real-time capability is crucial for

capturing sudden market shifts caused by global events, economic announcements, or company-specific developments.

f. Pre-processing of News Data

News classification using Natural Language Processing (NLP) is a key step in understanding the impact of textual information on stock market behavior. After collecting news articles, the text undergoes preprocessing steps such as tokenization, stop-word removal, stemming, and normalization to prepare it for analysis. NLP techniques are then applied to extract meaningful insights from the text. Sentiment analysis is performed to classify news as positive, negative, or neutral based on the tone and context of the content. This can be achieved using machine learning models, lexicon-based approaches, or deep learning techniques. The sentiment scores quantify investor perception and market psychology, which often influence stock price movements. Additionally, feature extraction methods such as TF-IDF or word embeddings may be used to represent text numerically. The classified sentiment data is then aligned with corresponding stock market timelines to ensure temporal relevance. By converting unstructured textual data into structured sentiment indicators, NLP enhances the model's ability to respond to real-world events. This integration of qualitative information with quantitative data improves prediction accuracy and provides a more holistic understanding of market dynamics.

IV. Result and discussion

The proposed hybrid framework combining Bidirectional Long Short-Term Memory (BiLSTM) and Bidirectional Gated Recurrent Unit (BiGRU) networks, along with NLP-based sentiment analysis, was evaluated using historical stock price data and real-time financial news. The model performance was assessed using standard metrics such as Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE). The experimental results demonstrate that the hybrid model achieves superior prediction accuracy compared to traditional machine learning approaches and individual deep learning models. The integration of BiLSTM and BiGRU enables the system to effectively capture both long-term dependencies and short-term fluctuations in stock price movements, resulting in improved learning of complex temporal patterns. Furthermore, the inclusion of technical indicators during feature extraction enhances the model's ability to identify

trends, momentum, and volatility within the data. A significant improvement in prediction performance was observed when sentiment analysis derived from financial news was incorporated. The NLP-based sentiment scores provided valuable insights into investor behavior and market psychology, particularly during periods of sudden market changes or high volatility. This demonstrates that combining quantitative data with qualitative information leads to more robust and reliable predictions. Comparative analysis reveals that the hybrid BiLSTM–BiGRU model consistently produces lower error values than standalone models, confirming the effectiveness of the proposed architecture. Additionally, the email notification system successfully generated timely alerts based on predicted trends, supporting real-world decision-making.

a. Input Layer

The input layer acts as the foundation of the BiGRU–BiLSTM model by receiving and organizing stock market data into a structured format suitable for sequential learning. The data consists of historical stock prices, technical indicators, and sentiment features, which are arranged as time-series sequences. Each input sample is represented as a sequence of feature vectors over a fixed number of time steps. The input is typically modeled as:

$$X = \{x_1, x_2, x_3, \dots, x_t\}$$

Where, x_t represents the feature vector at time step t .

This layer ensures that temporal relationships between data points are preserved, which is essential for capturing trends and patterns in stock price movements. Before feeding the data into the model, normalization or scaling techniques are applied to bring all feature values into a consistent range, improving training stability and convergence speed. Additionally, the input data is often transformed using sliding window techniques, where past observations are used to predict future values. The input layer itself does not perform computations or learning but plays a crucial role in structuring the data appropriately. Proper preparation at this stage ensures that the subsequent BiGRU and BiLSTM layers can effectively extract meaningful temporal features and dependencies from the data.

b. Bidirectional Recurrent Layer (BiGRU/BiLSTM Layer)

The bidirectional recurrent layer is the core component responsible for learning temporal dependencies in the data. It combines the strengths of BiLSTM and BiGRU to capture both long-term and short-term patterns in stock price sequences. The bidirectional mechanism processes the input sequence in both forward and backward directions, enabling the model to learn from past and future contexts simultaneously. The combined hidden representation is expressed as:

$$h_t = [\vec{h}_t, \overleftarrow{h}_t]$$

Here, $\vec{h}_t$ represents the forward hidden state and $\overleftarrow{h}_t$ represents the backward hidden state. This dual processing enhances the model's ability to understand complex dependencies and improves prediction accuracy. BiLSTM contributes by maintaining long-term memory through gated mechanisms, while BiGRU provides computational efficiency and faster learning of short-term dependencies. Together, they create a powerful hybrid architecture capable of modeling non-linear and dynamic stock market behavior. This layer extracts high-level features from sequential data, which are then passed to the next layer for final prediction.

c. Dense (Fully Connected) Output Layer

The dense layer is the final stage of the model, responsible for converting the extracted features into meaningful output predictions. It takes the high-level feature representation obtained from the bidirectional recurrent layers and applies a linear transformation to generate the predicted stock value. The operation of the dense layer can be represented as:

$$y = W \cdot h + b$$

Where, W is the weight matrix, h is the input feature vector from the previous layer, and b is the bias term.

This layer integrates all learned temporal features and produces a final output, such as the predicted stock price or trend. For regression tasks, a linear activation function is typically used to generate continuous values. The dense layer

plays a critical role in mapping complex learned patterns into actionable predictions. Its performance directly affects the overall accuracy of the model. Proper tuning of parameters such as the number of neurons and weights ensures better generalization and reduces prediction errors, making the model more reliable for real-time stock market forecasting.

d. Mean Squared Error (MSE)

Mean Squared Error (MSE) is a widely used evaluation metric for regression problems, particularly in stock price prediction models such as the proposed BiGRU–BiLSTM framework. It measures the average of the squared differences between the actual stock prices and the predicted values generated by the model. The formula for MSE is given by:

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

Where, y_i represents the actual stock price, $\hat{y}_i$ is the predicted value, and n is the number of observations. In the proposed system, MSE is used to evaluate how well the hybrid model captures complex market patterns. By squaring the errors, MSE penalizes larger deviations more heavily than smaller ones, making it particularly sensitive to outliers and large prediction mistakes. This is important in stock market forecasting, where significant prediction errors can lead to financial losses. A lower MSE value indicates that the model predictions are close to actual stock prices, reflecting better performance. During training, MSE is often used as a loss function to optimize the model parameters using gradient descent techniques. However, one limitation of MSE is that its squared nature may exaggerate the impact of extreme values. Despite this, it remains a fundamental metric for assessing model accuracy and guiding improvements in the proposed prediction system.

e. Root Mean Squared Error (RMSE)

Root Mean Squared Error (RMSE) is another important evaluation metric derived from MSE, commonly used to measure the accuracy of regression models like the proposed BiGRU–BiLSTM system. It represents the square root of the average squared differences between predicted and actual values. The formula for RMSE is:

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

RMSE provides an error measure in the same units as the original data, making it easier to interpret compared to MSE. In the context of stock market prediction, RMSE indicates the average magnitude of prediction error in terms of actual stock price values. A lower RMSE signifies higher prediction accuracy and better model performance. Similar to MSE, RMSE gives more weight to larger errors due to the squaring operation, which makes it sensitive to significant deviations in predictions. In the proposed system, RMSE is particularly useful for evaluating how well the model performs in real-world scenarios, where understanding the actual scale of prediction error is important for decision-making. Investors and analysts can use RMSE to assess the reliability of the model's forecasts. Although RMSE is sensitive to outliers, it provides a more interpretable measure of model performance, making it a preferred metric alongside MSE for evaluating the effectiveness of stock price prediction systems.

f. Mean Absolute Error (MAE)

Mean Absolute Error (MAE) is a simple and effective metric used to measure the average absolute difference between predicted and actual stock prices. The formula is:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

In the proposed BiGRU–BiLSTM system, MAE provides a clear and interpretable measure of prediction accuracy by treating all errors equally. Unlike MSE and RMSE, MAE does not square the errors, making it less sensitive to outliers and extreme values. This allows the model's performance to be evaluated based on typical prediction errors rather than being dominated by large deviations. A lower MAE indicates that the model consistently produces accurate predictions. MAE is particularly useful for understanding the average performance of the model under normal market conditions and is often used alongside MSE and RMSE for comprehensive evaluation.




Fig 4.1 for the loss graph for the proposed system

The loss graph of the proposed hybrid BiGRU–BiLSTM model illustrates the variation of training and validation loss across epochs, providing insight into the learning behavior of the system. Initially, both training and validation loss values decrease rapidly, indicating that the model is effectively learning underlying patterns in stock price data. As training progresses, the loss gradually stabilizes, showing convergence of the model. A small gap between training and validation loss suggests good generalization and minimal overfitting. The smooth decline in loss values confirms that the integration of NLP-based sentiment features and technical indicators enhances model performance. Overall, the loss graph demonstrates that the proposed system achieves stable training, reduced prediction error, and reliable forecasting capability.

V. Conclusion

In conclusion, this study presents an effective hybrid framework for stock market prediction by integrating Bidirectional Gated Recurrent Units (BiGRU) and Bidirectional Long Short-Term Memory (BiLSTM) networks with NLP-based sentiment analysis. The proposed model successfully captures both long-term and short-term dependencies in financial time-series data while incorporating external factors such as market sentiment derived from news. The bidirectional architecture enhances contextual learning, leading to improved prediction accuracy and reduced forecasting errors. Additionally, the inclusion of sentiment features strengthens the model's ability to adapt to dynamic market conditions and sudden fluctuations. Experimental results demonstrate that the hybrid approach outperforms traditional methods, providing more reliable and robust predictions. Future enhancements may include integrating advanced transformer models for better accuracy, incorporating social media sentiment analysis for deeper

market insights, and adding attention mechanisms to improve feature importance. The system can also be extended with reinforcement learning for automated trading and deployed as a real-time web or mobile application for better usability.

References

- [1] D. O. Oyewola and A. Abdulrahim, "Stock market prediction with optimized PLSTM-AL in smart urban cities," 2025.
- [2] Y. Chang, W. Lu, F. Xue, and X. Lu, "Combining market-guided patterns and Mamba for stock price prediction," 2025.
- [3] T. Uçkan, "Integrating PCA with deep learning models for stock market forecasting: An analysis of Turkish stock markets," 2024.
- [4] X. Chen, H. Xie, Z. Li, H. Zhang, X. Tao, and F. L. Wang, "Sentiment analysis for stock market research: A bibliometric study," 2025.
- [5] C. R. Kumar and S. Manikandan, "SEMP-TA: A novel stock market prediction approach based on stacking ensemble machine learning for effective trend analysis," 2025.
- [6] M. Si, Q. Li, and H. Li, "Stock movement prediction via attention-based LSTM and financial news sentiment," *IEEE Access*, 2020.
- [7] Z. Yang, Y. Yang, J. Wu, and Z. Tang, "Financial news sentiment analysis using BERT for stock trend prediction," *Expert Systems with Applications*, 2021.
- [8] Y. Li, K. Jiang, and J. Li, "Hybrid deep learning framework combining CNN-LSTM and technical indicators for stock price forecasting," *Knowledge-Based Systems*, 2021.
- [9] R. Sawhney, A. Agarwal, and R. Shah, "Stock price prediction using FinBERT and deep learning models," *Proceedings of the ACL Financial NLP Workshop*, 2022.
- [10] H. Chen and W. Qian, "Multimodal stock prediction using price series and news sentiment with transformers," *Information Processing & Management*, 2022.
- [11] S. Kim and H. Kim, "Temporal convolutional network based financial time-series forecasting," *IEEE Access*, 2023.



[12] L. Zhang, P. Zhao, and Y. Chen, "Fusion of technical indicators and BERT sentiment embeddings for stock market prediction," *Applied Soft Computing*, 2023. [13] K. Patel and M. Shah, "Explainable AI-driven stock trend prediction using transformer architectures," *Expert Systems with Applications*, 2024.

[14] J. Wang, X. Liu, and Y. Sun, "Real-time stock forecasting using hybrid TCN and attention-based sentiment analysis," *Knowledge-Based Systems*, 2024.

[15] R. Gupta and S. Verma, "Adaptive multimodal deep learning framework for financial market prediction using news and price dynamics," *IEEE Transactions on Artificial Intelligence*, 2025