

Stock Price Prediction Using Deep Learning Models: A Comparative Study of LSTM, GRU, Bi-LSTM, and Hybrid Architectures

Sayeed Baig¹, Venkat Baba Yemineni², Vinayak Arun Naik³, Sameeksha Uday Naik⁴, Shine Malik⁵, Dr. Savitha G⁶

^{1,2,3,4,5}Department of Computer Science and Engineering, RV Institute of Technology and Management,

⁶Professor, Department of Computer Science and Engineering, RV Institute of Technology and Management, Bengaluru, Karnataka, India

Abstract - Stock price prediction is a challenging task due to the volatile and non-linear nature of financial markets. This study explores the effectiveness of deep learning models for time-series forecasting of stock prices using historical data from Tata Consultancy Services (TCS). Multiple models, including Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), Bidirectional LSTM (Bi-LSTM), and a hybrid LSTM-GRU architecture, were implemented and evaluated. The dataset was preprocessed using normalization techniques and structured into sequential time windows to capture temporal dependencies. Each model was trained using historical closing prices, and performance was evaluated using metrics such as Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), R^2 score, and classification-based metrics for directional accuracy. Experimental results indicate that the GRU model achieved the best performance with the lowest prediction error (RMSE \approx 42.50 and MAE \approx 30.66) and highest R^2 score (0.9358), demonstrating superior capability in capturing time-series patterns. The hybrid LSTM-GRU model also showed strong performance, while LSTM and Bi-LSTM exhibited comparatively lower accuracy. These findings highlight the effectiveness of GRU-based architectures for stock price prediction tasks. This study provides a comparative framework for evaluating deep learning models in financial forecasting and offers insights into selecting appropriate architectures for real-world stock prediction applications.

Key Words: Stock Price Prediction, Deep Learning, LSTM, GRU, Bi-LSTM, Time Series Forecasting, Financial Data Analysis

1. INTRODUCTION

The prediction of stock prices has long been a significant challenge in the field of financial analysis due to the highly volatile and non-linear nature of stock markets. Accurate forecasting of stock prices can assist investors, financial analysts, and organizations in making informed decisions, thereby reducing risks and maximizing profits. Traditional statistical methods often fail to capture complex temporal patterns present in financial data, leading to the exploration of advanced machine learning and deep learning techniques. In recent years, deep learning models have shown promising results in time-series forecasting tasks due to their ability to learn long-term dependencies and non-linear relationships. Models such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) are specifically designed to handle sequential data and have been widely applied in financial forecasting. Additionally, hybrid and bidirectional architectures further enhance the learning capability by capturing more complex patterns in data.

This study focuses on predicting stock prices using multiple deep learning models, including LSTM, GRU, Bi-LSTM, and a hybrid LSTM-GRU model. The objective is not only to forecast stock prices but also to perform a comparative analysis of these models to determine the most effective architecture for time-series prediction. The models are trained on historical stock data of Tata Consultancy Services (TCS), and their performance is evaluated using various statistical metrics.

The remainder of this paper is organized as follows: Section II discusses related work, Section III presents the methodology, Section IV describes experimental results and analysis, and Section V concludes the study with future directions.

2. LITERATURE REVIEW

Stock price prediction has been extensively studied using various statistical, machine learning, and deep learning approaches. Traditional methods such as Autoregressive Integrated Moving Average (ARIMA) and linear regression have been widely used for time-series forecasting. However, these methods often struggle to capture non-linear patterns and complex dependencies present in financial data. With the advancement of machine learning, models such as Support Vector Machines (SVM), Random Forest, and Artificial Neural Networks (ANN) have been applied to stock prediction problems. While these models improved prediction performance compared to traditional techniques, they still have

limitations in handling sequential dependencies effectively. Deep learning models, particularly recurrent neural networks (RNNs), have shown significant improvements in time-series forecasting tasks. Among them, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models are widely used due to their ability to retain long-term dependencies and mitigate the vanishing gradient problem. Furthermore, advanced architectures such as Bidirectional LSTM (Bi-LSTM) and hybrid models combining multiple networks have been proposed to enhance prediction accuracy by capturing both past and future context.

Despite these advancements, selecting the most effective model for stock price prediction remains a challenge. Therefore, this study focuses on a comparative analysis of multiple deep learning models to evaluate their performance on real-world stock data.

3. METHODOLOGY

This section describes the dataset, preprocessing techniques, and deep learning models used for stock price prediction. The overall workflow includes data collection, preprocessing, sequence generation, model training, and evaluation.

3.1 Dataset Description

The dataset used in this study consists of historical stock data of Tata Consultancy Services (TCS). It includes features such as Date, Open, High, Low, Close, and Volume. Among these, the closing price is primarily used for prediction as it represents the final market value of the stock for a given day.

The dataset spans multiple years, providing sufficient data for training and testing the deep learning models. The data is first sorted chronologically to preserve temporal relationships essential for time-series forecasting.

3.2 Data Preprocessing

Data preprocessing is a crucial step to ensure effective model performance. Initially, irrelevant or missing values are handled, and the dataset is structured appropriately. The closing price values are extracted and normalized using MinMaxScaler to scale the data between 0 and 1, which helps improve the learning efficiency of neural networks.

To convert the data into a time-series format, a sliding window approach is used. In this study, a time step of 60 days is considered, meaning that the model uses the past 60 days of stock prices to predict the next day's price. The dataset is then split into training and testing sets to evaluate model performance.

3.3 Model Architectures

Multiple deep learning models are implemented in this study to analyze their effectiveness in stock price prediction. The models used include Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), Bidirectional LSTM (Bi-LSTM), and a hybrid LSTM-GRU model.

The LSTM model is capable of learning long-term dependencies in sequential data, making it suitable for time-series forecasting. The GRU model is a simplified version of LSTM that offers faster training while maintaining comparable performance. The Bi-LSTM model processes data in both forward and backward directions to capture additional context. The hybrid LSTM-GRU model combines the strengths of both architectures to improve predictive performance. Each model is trained using sequential input data and optimized using appropriate loss functions and optimizers.

3.4 Model Architecture Diagrams

The architecture of each deep learning model used in this study is illustrated to provide a clear understanding of the data flow and internal structure. These diagrams represent how sequential input data is processed through different layers to generate predictions.

LSTM Architecture

The LSTM model consists of sequential layers designed to capture long-term dependencies in time-series data. The input sequence, consisting of past 60 time steps, is fed into an LSTM layer with multiple hidden units. The LSTM layer processes the sequence through memory cells and gating mechanisms, enabling it to retain relevant information over time. The output of the LSTM layer is then passed to a fully connected dense layer, which produces the final predicted stock price.

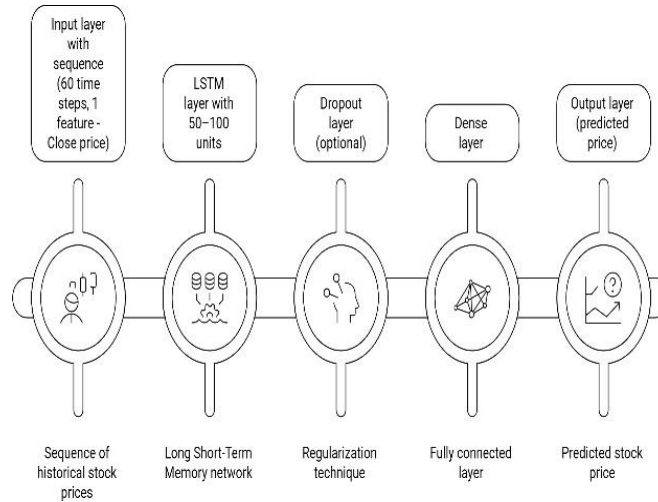


Chart -1: Architecture of LSTM Model

GRU Architecture

The GRU model follows a structure similar to LSTM but with a simplified architecture. It consists of a GRU layer that processes sequential input data using update and reset gates. This reduces computational complexity while maintaining the ability to capture temporal dependencies. The output from the GRU layer is passed through a dense layer to generate the predicted value. The GRU model demonstrates efficient learning with fewer parameters compared to LSTM.

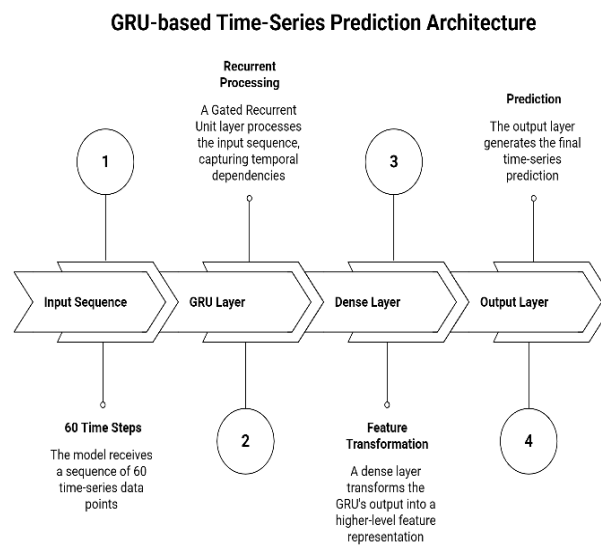


Fig -1: Architecture of GRU Model

Bi-LSTM Architecture

The Bidirectional LSTM (Bi-LSTM) model processes the input sequence in both forward and backward directions. This allows the model to capture past and future context simultaneously. The outputs from both directions are combined and passed to a dense layer for final prediction. This architecture enhances the model’s ability to understand complex patterns in sequential data.

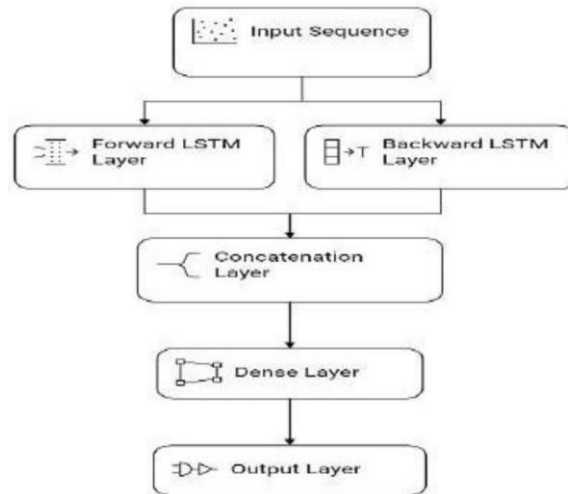


Fig -2: Architecture of Bi-LSTM Model

Hybrid LSTM-GRU Architecture

The hybrid model combines LSTM and GRU layers to leverage the advantages of both architectures. The input sequence is first processed by LSTM layers to capture longterm dependencies. The output is then passed to GRU layers, which further refine the learned features with reduced complexity. Additionally, normalization and attention mechanisms are applied to improve feature representation. Finally, fully connected layers are used to produce the output. This hybrid approach aims to improve prediction accuracy by combining complementary strengths of both models.

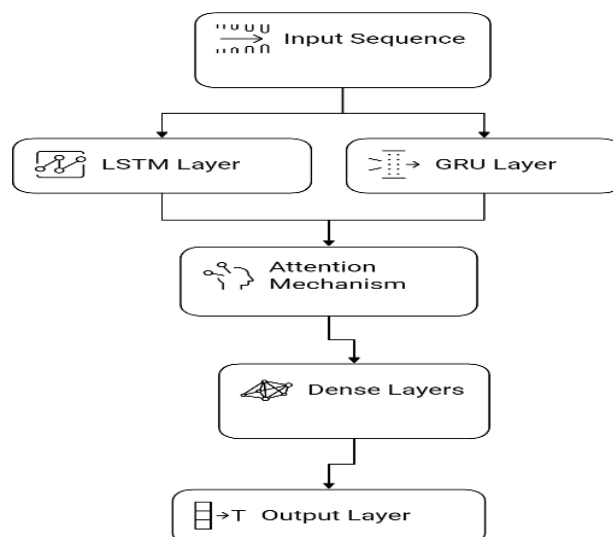


Fig -3: Architecture of Hybrid LSTM-GRU Model

3.5 Training Process

The dataset is divided into training and testing sets, where approximately 80% of the data is used for training and the remaining 20% for testing. The models are trained over multiple epochs, allowing them to iteratively learn patterns in the data and minimize prediction error.

During training, the models adjust their internal parameters based on the loss value, which represents the difference between predicted and actual values. Techniques such as early stopping are applied in certain models to prevent overfitting and improve generalization performance.

3.6 Evaluation Metrics and Mathematical Formulation

The performance of the proposed models is evaluated using standard regression metrics such as Mean Absolute Error (MAE) and Mean Squared Error (MSE). These metrics provide a quantitative measure of the difference between predicted and actual stock prices.

Mean Absolute Error (MAE)

Mean Absolute Error (MAE) measures the average magnitude of errors in predictions, without considering their direction. It provides a straightforward interpretation of prediction accuracy

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

Mean Squared Error (MSE)

Mean Squared Error (MSE) measures the average of the squared differences between predicted and actual values. It penalizes larger errors more significantly, making it useful for emphasizing large deviations.

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

4. RESULTS AND DISCUSSION

This section presents the experimental results obtained from different deep learning models and provides a comparative analysis of their performance. The models were evaluated using various statistical metrics and visual analysis of predicted versus actual stock prices.

The LSTM model demonstrated moderate performance with higher error values and lower predictive accuracy compared to other models. The directional accuracy was also relatively low, indicating limitations in capturing market trends effectively.

The GRU model outperformed other models by achieving the lowest error values and the highest R^2 score, indicating strong predictive capability and better generalization on unseen data. Its simpler architecture allowed efficient learning of time-series patterns.

The hybrid LSTM-GRU model showed competitive performance by combining the strengths of both architectures. Although it performed well, it did not surpass the GRU model in terms of overall accuracy.

The Bi-LSTM model provided moderate classification-based performance, with an accuracy of approximately 54%, indicating its ability to capture directional trends but with limited precision.

Overall, the comparative analysis indicates that the GRU model is the most effective among all the implemented models for stock price prediction in this study.



Chart -1: Actual vs Predicted Stock Prices using LSTM Model

The LSTM model shows a reasonable trend alignment between actual and predicted stock prices. However, noticeable deviations can be observed during sudden market fluctuations, indicating limitations in capturing short-term volatility

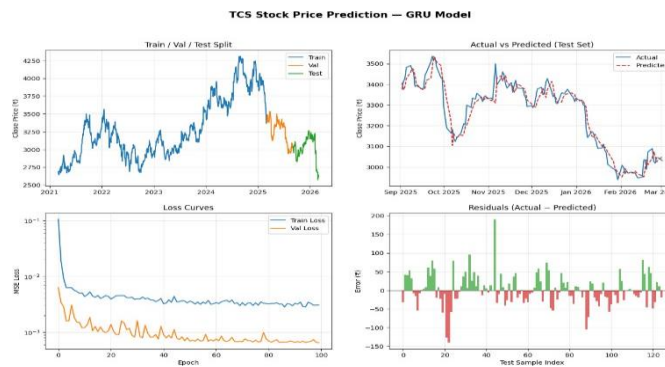


Chart -2: Actual vs Predicted Stock Prices using GRU Model

The GRU model demonstrates a closer alignment between actual and predicted values compared to LSTM. The model effectively captures both upward and downward trends, resulting in improved prediction accuracy and reduced error.

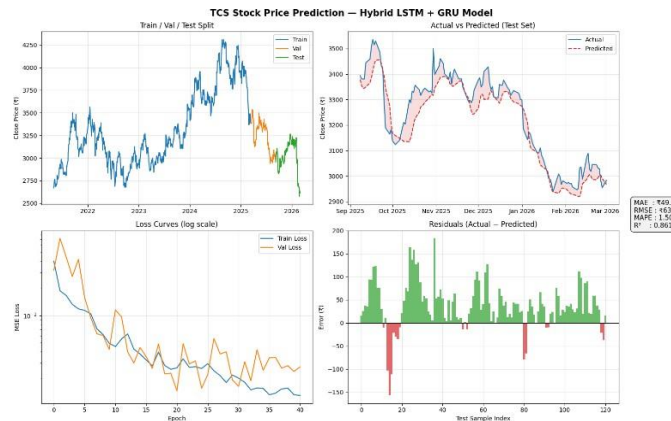


Chart -3: Actual vs Predicted Stock Prices using Hybrid LSTMGRU Model

The hybrid LSTM-GRU model combines the strengths of both architectures and produces stable predictions. While it performs better than the standalone LSTM model, slight deviations still exist when compared to the GRU model.

Table -1: Sample Table format

Model	MAE %	RMSE %	R ² Score	Accuracy/MAPE
LSTM	69.56	92.23	0.5"	48.5%
GRU	30.66	42.50		99.06%
Bi-LSTM	—	—	—	54.3%
GRU-LSTM	49.07	63.24	0.8618	98.5%

From Table I, it is evident that the GRU model achieves the best performance with the lowest error values and highest R² score. The hybrid model also performs well but does not outperform GRU. The LSTM model shows comparatively higher error, while Bi-LSTM provides moderate directional prediction capability

5. CONCLUSION AND FUTURE WORK

In this study, multiple deep learning models, including LSTM, GRU, Bi-LSTM, and a hybrid LSTM-GRU architecture, were implemented to predict stock prices using historical time-series data of Tata Consultancy Services (TCS). The models were evaluated using various performance metrics such as MAE, RMSE, MAPE, and R² score, along with graphical analysis of predicted and actual values.

The experimental results indicate that the GRU model outperforms other models in terms of prediction accuracy and error minimization. The hybrid LSTM-GRU model also demonstrates strong performance, while the LSTM and BiLSTM models show comparatively lower effectiveness in capturing complex market patterns. These findings highlight the importance of selecting appropriate deep learning architectures for financial time-series forecasting.

Despite achieving promising results, stock price prediction remains a challenging task due to market volatility and external influencing factors that are not captured in historical data alone. Therefore, future work can focus on incorporating additional features such as news sentiment analysis, macroeconomic indicators, and real-time data to further improve prediction accuracy. Additionally, advanced architectures such as attention-based models and transformer networks can be explored for better performance.

REFERENCES

- [1] S. Kalange et al., "Stock Price Prediction Using Machine Learning and Deep Learning Techniques," AJPAS, 2025.
- [2] "Stock Price Prediction Using Deep Learning Models," arXiv:2505.05325v1, 2025.
- [3] "A Hybrid LSTM-GRU Model for Stock Price Prediction," 2024.
- [4] Z. Chen et al., "Stock market prediction using deep neural networks," IEEE Access, 2022.
- [5] S. Basak et al., "Predicting stock market price using hybrid models," IEEE, 2022.
- [6] T. Fischer et al., "Deep learning with LSTM for stock market prediction," 2021 update.
- [7] J. Bao et al., "A deep learning framework for financial time series using stacked LSTM," IEEE Access, 2021.
- [8] Y. Wang et al., "A hybrid LSTM model for stock prediction," IEEE Access, 2020.
- [9] H. Sezer et al., "Financial time series forecasting with deep learning: A systematic review," Applied Soft Computing, 2020.
- [10] S. Selvin et al., "Stock price prediction using LSTM, RNN and CNN sliding window model," Procedia Computer Science, 2020.