Comparative Performance of LSTM and GRU in Forecasting NVIDIA Stock Prices: Insights from the Digital Financial Market

Richard Berlien[1,*]

[1]Universitas Murni Teguh, Indonesia

ABSTRACT

The rapid evolution of artificial intelligence and digital technologies has transformed financial forecasting, enabling more accurate and data-driven decision-making in digital markets. This study investigates the comparative performance of two deep learning architectures, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), in predicting NVIDIA stock prices as a case study of digital financial market behavior. Historical daily stock data from 2018 to 2026 were analyzed using key financial indicators, including opening, closing, high, low, and trading volume. Both models were developed using a time-series forecasting framework with data normalization, a 60-day sliding window, and evaluation through Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). The results show that both LSTM and GRU models effectively captured temporal dependencies and nonlinear price movements, with the GRU model achieving superior accuracy (RMSE 3.5801, MAE 2.4518) and faster convergence compared to LSTM (RMSE 4.5841, MAE 3.1153). The GRU model also demonstrated higher adaptability to short-term fluctuations, while LSTM exhibited stronger long-term trend stability. These findings confirm that deep learning architectures can significantly enhance stock price prediction accuracy and offer valuable insights for digital market analysis. The study contributes to the growing application of AI in finance by highlighting the potential of GRU-based forecasting models for real-time prediction and decision support in dynamic digital trading environments.

Keywords Deep Learning, LSTM, GRU, Stock Price Prediction, Digital Financial Market

INTRODUCTION

The advancement of artificial intelligence and digital technology has transformed the global financial landscape, particularly in market forecasting and investment decision-making [1]. As stock markets become increasingly complex and volatile, investors and analysts are relying more heavily on intelligent systems to generate data-driven insights. Traditional statistical methods, such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH), have been widely employed for stock price prediction [2]. However, these models have limitations in capturing nonlinear patterns and the highly dynamic characteristics of digital financial markets. As trading volumes and data sources continue to expand, conventional forecasting approaches face challenges in processing large-scale sequential data effectively. These limitations have encouraged the adoption of deep learning techniques, which can capture complex temporal dependencies and potentially improve predictive accuracy in financial systems [3].

In recent years, deep learning models have emerged as promising approaches for time-series prediction because of their ability to model nonlinear and non-

Submitted 1 May 2026
Accepted 15 July 2026
Published 27 October 2026

Corresponding author
Richard Berlien,
richardberlien@gmail.com

Additional Information and Declarations can be found on page 291

DOI: 10.63913/ftij.v2i4.40

Copyright
2026 Berlien

Distributed under
Creative Commons CC-BY 4.0

How to cite this article: R. Berlien, "Comparative Performance of LSTM and GRU in Forecasting NVIDIA Stock Prices: Insights from the Digital Financial Market," *FinTech Innov. J.*, vol. 2, no. 4, pp. 280-293, 2026.

stationary data. Among these models, recurrent neural networks (RNNs) and their advanced variants, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), have attracted considerable attention in financial forecasting research. LSTM and GRU architectures incorporate gating mechanisms that help regulate information flow and mitigate the vanishing gradient problem commonly encountered in traditional RNNs. Previous research has demonstrated that LSTM models can outperform conventional machine learning approaches, including Random Forest and Support Vector Machines (SVM), in capturing long-term dependencies in financial data [4]. Other studies have reported that GRU models can achieve predictive performance comparable to LSTM while offering greater computational efficiency due to their simpler architecture [5]. These characteristics make LSTM and GRU relevant approaches for financial time-series prediction, particularly in market environments characterized by frequent fluctuations and complex temporal patterns.

Despite these advancements, several research gaps remain in the existing literature. Many studies have evaluated LSTM or GRU independently rather than directly comparing their performance under identical experimental conditions. Furthermore, existing research has often focused on short-term datasets or specific market segments, potentially limiting the generalizability of its findings. Comprehensive evaluations using extended historical datasets and multiple input features, including opening, closing, high, low, and trading volume, remain important for understanding the predictive capabilities of these architectures. These variables provide complementary information about market behavior but may not always be incorporated comprehensively in forecasting models. In addition, further investigation is needed to determine how effectively LSTM and GRU models capture the price dynamics of technology-driven assets such as NVIDIA, whose stock performance is influenced by developments in the semiconductor industry, artificial intelligence, and the broader digital economy.

This study addresses these research gaps by conducting a comparative analysis of LSTM and GRU models for NVIDIA stock price forecasting using historical data spanning 2018 to 2026. NVIDIA was selected as the case study because of its significant role in the semiconductor and artificial intelligence industries and the dynamic nature of its stock price movements. The research adopts a structured experimental design involving data normalization, time-series windowing, and model evaluation using Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). By applying consistent preprocessing procedures and comparable model configurations, the study establishes a controlled basis for evaluating the forecasting performance of both architectures. This approach enables an objective assessment of their predictive capabilities and provides empirical evidence regarding their relative strengths and limitations in financial time-series forecasting.

The contribution of this research extends beyond a direct comparison of deep learning architectures. By evaluating the predictive accuracy of LSTM and GRU models using multi-year historical data, this study provides empirical insights into the application of artificial intelligence in stock market forecasting. The findings can help clarify the relative suitability of these architectures for financial prediction tasks and inform the development of AI-driven financial analytics and

decision-support systems. The results may also provide useful information for investors, portfolio managers, and financial institutions seeking to incorporate predictive models into investment analysis and risk management. Furthermore, this study contributes to the broader understanding of deep learning performance in financial contexts and supports further research into intelligent forecasting systems for digital financial markets.

Literature Review and Related Works

Stock price forecasting has been a long-standing research topic due to its critical role in financial decision-making, risk management, and investment strategy formulation. Early studies in this domain primarily employed traditional statistical and econometric models such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) to model historical price movements [6], [7]. Although these models provide strong mathematical foundations, they rely heavily on assumptions of linearity and stationarity, which often fail to represent the complex and nonlinear behavior of real financial markets. As a result, their forecasting accuracy tends to decline in highly volatile and dynamic market environments [8].

To address these limitations, machine learning techniques were introduced to enhance predictive performance by capturing nonlinear relationships in financial time-series data. Models such as Support Vector Machines, Random Forests, and feedforward Artificial Neural Networks demonstrated improved accuracy compared to traditional statistical approaches [9], [10]. However, these models typically treat observations as independent samples and do not explicitly account for temporal dependencies across time. This limitation reduces their effectiveness when applied to sequential financial data, where past information plays a crucial role in determining future price movements [11].

The emergence of deep learning marked a significant shift in time-series forecasting, particularly with the introduction of Recurrent Neural Networks (RNNs) [12]. RNNs are designed to process sequential data by maintaining internal states that capture temporal information. Despite their potential, standard RNNs suffer from the vanishing gradient problem, which restricts their ability to learn long-term dependencies in extended sequences [13]. To overcome this challenge, advanced RNN architectures such as Long Short-Term Memory (LSTM) were developed, incorporating memory cells and gating mechanisms that enable the preservation of relevant information over long time horizons [14]. Numerous studies have demonstrated that LSTM-based models outperform traditional machine learning and statistical models in stock price prediction tasks [15], [16].

Gated Recurrent Unit (GRU) was later introduced as a simplified alternative to LSTM, reducing architectural complexity while maintaining strong performance in sequential learning tasks [17]. By combining gating mechanisms into a more compact structure, GRU reduces computational cost and training time. Several studies have reported that GRU achieves comparable or superior forecasting accuracy relative to LSTM, particularly in scenarios involving short-term prediction and highly volatile data [18]. These characteristics make GRU especially attractive for real-time financial forecasting applications, where efficiency and responsiveness are essential.

Despite the growing body of research on deep learning-based stock prediction, there remains a lack of comprehensive comparative studies that evaluate LSTM and GRU under identical experimental settings using long-term datasets and consistent evaluation metrics. Many existing studies focus on limited time spans or specific market conditions, which restricts the generalizability of their findings [19]. Furthermore, relatively little attention has been given to technology-driven stocks that represent the digital economy, such as NVIDIA, whose price behavior is closely linked to innovation cycles and global technological trends [20]. Addressing these gaps, the present study conducts a systematic comparison of LSTM and GRU models using multi-year NVIDIA stock data, thereby contributing empirical insights to the advancement of deep learning applications in digital financial market forecasting.

Methodology

This study employs a quantitative experimental approach to compare the forecasting performance of Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models in predicting NVIDIA stock prices. The overall research procedure follows a structured and sequential framework, as illustrated in figure 1, which outlines the main stages of the study. The research process begins with data collection of historical NVIDIA stock prices, followed by data preprocessing and normalization to ensure numerical stability.

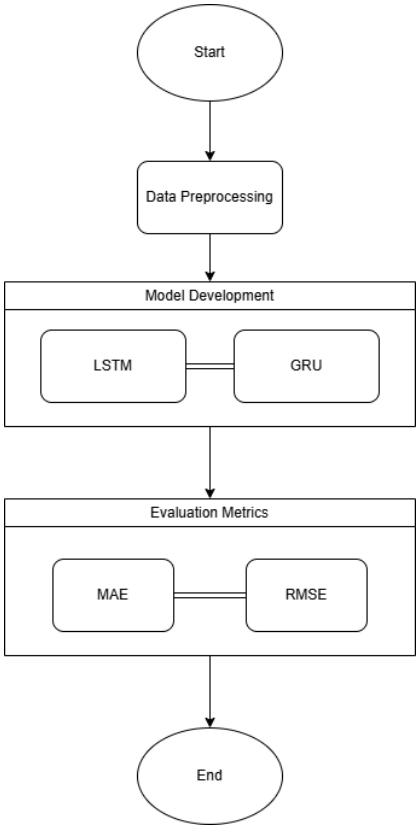


Figure 1 Research Steps

The preprocessed data are then transformed into a supervised time-series format using a sliding window approach. Subsequently, two deep learning models, LSTM and GRU, are developed and trained using identical

configurations to ensure a fair comparison. The trained models are evaluated using statistical error metrics, and the results are analyzed through quantitative measurements and visual interpretations. This systematic workflow ensures methodological consistency, reproducibility, and objective performance assessment in the context of digital financial market forecasting.

To improve training stability and accelerate convergence, the closing price data were normalized using Min–Max scaling, which transforms the original values into a bounded range between 0 and 1. The normalization process is defined as:

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}} X \tag{1}$$

represents the original closing price, and $X_{min}$and X_{max}denote the minimum and maximum values of the dataset, respectively. This scaling technique preserves the relative structure of the data while preventing larger values from dominating the learning process. After normalization, a sliding window approach was applied to convert the time-series data into a supervised learning format. A window size of 60 trading days was used, meaning that the closing prices from the previous 60 days were employed as input features to predict the closing price of the next day. This transformation enables the models to learn temporal dependencies and historical patterns embedded in the financial time series.

The transformed dataset was then split sequentially into training and testing sets using an 80:20 ratio to preserve the chronological order of the data and avoid information leakage. The training set was used to optimize model parameters, while the testing set was reserved for performance evaluation. Both LSTM and GRU models were designed using identical architectures to ensure a fair comparison. Each model consists of two recurrent layers with 50 hidden units, followed by dropout layers with a dropout rate of 0.2 to reduce overfitting. The final output layer is a fully connected dense layer with a single neuron that produces the predicted closing price.

The LSTM model captures long-term dependencies through its memory cell and gating mechanisms, which regulate information flow using input, forget, and output gates. The internal operations of an LSTM cell can be expressed as:

$$\begin{aligned} f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \tilde{C}_t \\ = tanh(W_C \cdot [h_{t-1}, x_t] + b_C) C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t o_t \\ = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) h_t = o_t \cdot tanh(C_t) x_t \end{aligned} \tag{2}$$

is the input at time step t, $h_t$is the hidden state, $C_t$is the cell state, and $\sigma(\cdot)$denotes the sigmoid activation function. In contrast, the GRU model simplifies the gating structure by combining memory mechanisms into update and reset gates, which reduces computational complexity. The GRU operations are defined as:

$$\begin{aligned} z_t = \sigma(W_z \cdot [h_{t-1}, x_t] + b_z) r_t = \sigma(W_r \cdot [h_{t-1}, x_t] + b_r) \tilde{h}_t \\ = tanh(W_h \cdot [r_t \cdot h_{t-1}, x_t] + b_h) h_t \\ = (1 - z_t) \cdot h_{t-1} + z_t \cdot \tilde{h}_t \end{aligned} \tag{3}$$

These formulations allow GRU to retain relevant historical information while adapting more quickly to new inputs.

Model performance was evaluated using Root Mean Square Error (RMSE) and

Mean Absolute Error (MAE), which measure prediction accuracy and error magnitude. These metrics are computed as:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}\ ;\ MAE = \frac{1}{n}\sum_{i=1}^{n} |\ y_i - \hat{y}_i\ | \quad (4)$$

Lower RMSE and MAE values indicate better predictive performance. In addition to numerical evaluation, several visualization techniques were employed, including predicted versus actual price plots, training loss curves, residual error plots, correlation heatmaps, and distribution comparisons. These visual analyses provide deeper insights into model behavior, convergence characteristics, and generalization capability, thereby supporting a comprehensive assessment of deep learning-based forecasting in digital financial markets.

Algorithm 1 LSTM–GRU Stock Price Forecasting

Input:
Historical stock price data D, window size w, training ratio α

Output:
Predicted stock prices, RMSE, MAE

Process:

Start

Load and clean historical stock price data D

Normalize closing price using Min–Max normalization

$$x_i^{norm} = \frac{x_i - x_{min}}{x_{max} - x_{min}}$$

Construct time-series data using sliding window

$$X_i = \{x_i^{norm}, x_{i+1}^{norm}, \dots, x_{i+w-1}^{norm}\}, y_i = x_{i+w}^{norm}$$

Split dataset into training and testing sets

Initialize LSTM and GRU models with identical architecture

Train both models using training data

Predict stock prices using testing data

Apply inverse normalization to predictions

Evaluate model performance

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

Compare LSTM and GRU performance

End

Results

The experimental results provide a detailed comparison of the predictive performance of the Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models in forecasting NVIDIA stock prices. Both models were evaluated using two statistical performance metrics, Root Mean Square Error (RMSE) and Mean Absolute Error (MAE), to determine their forecasting accuracy and error magnitude. The results show that both architectures are capable of learning the temporal dependencies inherent in financial time-series data, yet their performance differs in precision and efficiency. As shown in table 1, the GRU

model achieved an RMSE value of 3.5801 and an MAE of 2.4518, while the LSTM model recorded an RMSE of 4.5841 and an MAE of 3.1153. These results suggest that GRU provides more accurate predictions and demonstrates a better fit to the actual data patterns compared to LSTM.

The lower error values obtained by the GRU model indicate its ability to generalize effectively across the testing dataset while maintaining computational efficiency. This improvement can be attributed to the simplified architecture of GRU, which reduces the number of parameters and training time without compromising its capacity to capture sequential dependencies. The LSTM model, although slightly less accurate, still performed competitively and produced stable results throughout the testing phase. The difference in RMSE and MAE between the two models reflects the distinct internal mechanisms of their gating structures, where GRU's update and reset gates enable faster adaptation to fluctuations in stock price movements. Overall, the quantitative findings confirm that GRU delivers a more efficient and slightly more accurate solution for short-term financial forecasting compared to LSTM.

Table 1 Comparative performance of LSTM and GRU models in NVIDIA stock price forecasting.

Model	RMSE	MAE
LSTM	4.5841	3.1153
GRU	3.5801	2.4518

Figure 2 presents a detailed comparison between the actual NVIDIA stock prices and the predictions generated by the LSTM and GRU models. Both models were able to replicate the general upward trend of the stock prices, indicating that each architecture successfully captured the temporal dependencies within the financial data. The predicted curves for LSTM and GRU followed the movement of the actual prices with minimal deviation, particularly during periods of moderate volatility. This consistency demonstrates the capability of both deep learning models to recognize patterns and relationships across sequential data points. However, subtle differences are noticeable in how each model responds to fluctuations in the market. The GRU model showed a stronger ability to adjust to changes in price movements, producing predictions that more closely matched the peaks and troughs of the actual price trend.

In contrast, the LSTM model generated smoother prediction curves, reflecting its strength in learning long-term dependencies within the data. This smoothness can be advantageous for capturing general market trends over extended time periods but may cause a delay in responding to rapid or short-term price changes. The GRU's greater responsiveness suggests that it can adapt more quickly to sudden variations, making it more effective for forecasting in highly dynamic trading environments. Despite these differences, both models maintained high predictive reliability, with their outputs exhibiting only minor deviations from real stock price movements. The visual alignment between the predicted and actual prices supports the quantitative results, confirming that both LSTM and GRU are capable of modeling complex, non-linear behaviors in stock market time-series data.

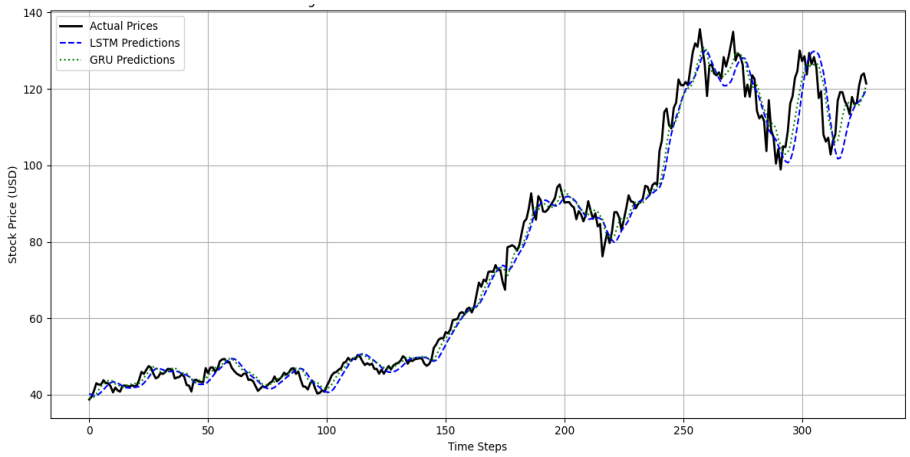


Figure 2 LSTM vs GRU prediction on NVIDIA stock prices.

Figure 3 illustrates the training performance of both the LSTM and GRU models throughout the learning process. At the beginning of training, both models experienced a sharp decline in loss values, indicating that they were quickly learning the underlying patterns in the stock price data. This rapid reduction in loss during the early epochs shows that the models were able to efficiently capture the main trends in the dataset and adjust their internal weights accordingly. After approximately the tenth epoch, the loss values for both models began to stabilize, suggesting that they had reached a state of convergence and were no longer making large adjustments to their parameters. The stability of the loss curve beyond this point reflects that both models achieved a balanced learning rate, avoiding overfitting or underfitting during training.

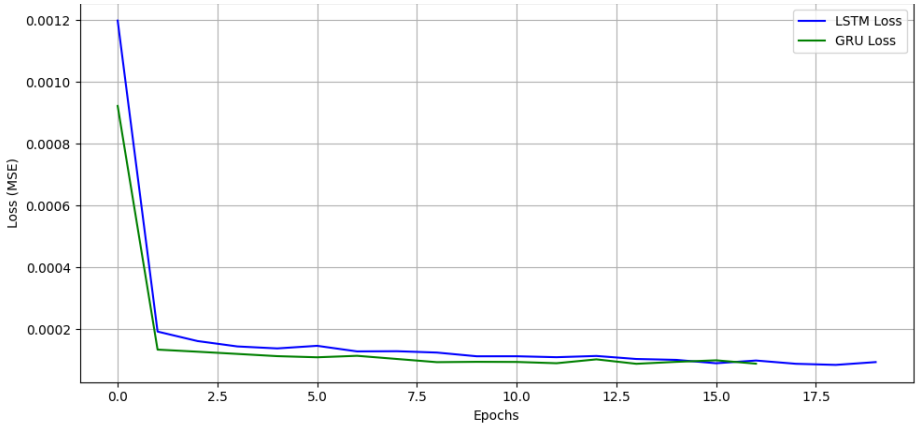


Figure 3 Training loss comparison between LSTM and GRU models.

Throughout the training phase, the GRU model consistently maintained lower loss values compared to the LSTM model. This outcome indicates that GRU reached convergence more efficiently and required fewer computational resources to achieve optimal performance. The improved optimization efficiency of GRU can be attributed to its simpler internal structure, which enables faster information propagation and gradient updates. In contrast, the LSTM model, which includes an additional cell state component, demonstrated a slower decline in loss but maintained smooth learning behavior. Despite this difference,

both models achieved effective training stability, with GRU showing a slight advantage in convergence speed. These results reinforce the notion that GRU can deliver high predictive accuracy while offering computational advantages for large-scale or time-sensitive forecasting tasks in financial modeling.

Figure 4 presents the residual error analysis for both the LSTM and GRU models, illustrating the difference between the actual and predicted NVIDIA stock prices over the testing period. The residual plots show that the errors for both models are evenly distributed around zero, which indicates that the models did not consistently overestimate or underestimate the stock prices. This symmetric distribution confirms that both LSTM and GRU produced unbiased predictions and were able to generalize effectively to unseen data. Most of the residuals remain within a narrow range, signifying that the magnitude of prediction errors was relatively small throughout the test set. This outcome demonstrates the capability of both models to maintain accuracy even during volatile market conditions, where rapid fluctuations in stock prices can often lead to larger prediction errors in traditional forecasting methods.

Although both models achieved balanced error distributions, the GRU model produced residuals with slightly smaller amplitude compared to the LSTM model. This pattern suggests that GRU predictions were more consistent, with fewer extreme deviations from the actual values. The reduced variance in GRU residuals highlights its ability to adapt more efficiently to short-term changes in the data while maintaining prediction stability. In contrast, the LSTM residuals exhibited slightly wider dispersion, reflecting minor delays in capturing sudden market movements due to its stronger emphasis on long-term dependencies. Overall, the residual analysis reinforces the quantitative results reported earlier, where GRU achieved lower RMSE and MAE values. These findings confirm that GRU provided more precise and stable predictions, making it particularly suitable for applications that require high accuracy in short-term financial forecasting.

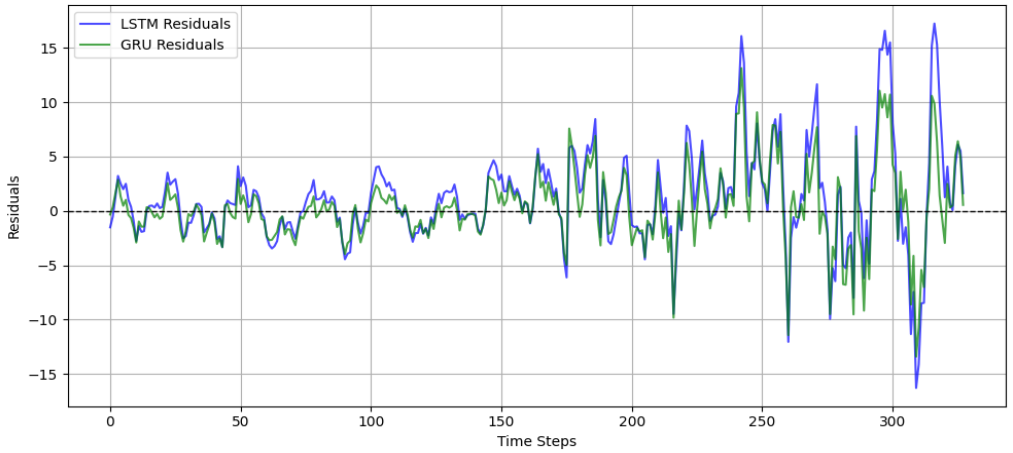


Figure 4 Residual error plot for LSTM and GRU models (Actual – Predicted).

Figure 5 illustrates the comparison between the actual stock price distribution and the predicted price distributions generated by the LSTM and GRU models. The overlapping curves demonstrate that both models were able to capture the underlying statistical patterns of NVIDIA's stock price movements with high fidelity. The actual price distribution exhibits a slightly right-skewed pattern,

which is typical of financial market data where upward price movements are often followed by smaller corrective declines. Both LSTM and GRU models successfully reproduced this distribution shape, indicating that their predictions align closely with the real-world variability of the stock prices. This similarity between the predicted and actual distributions highlights the ability of deep learning models to learn complex probabilistic relationships embedded within sequential financial data. It also shows that the models are not only accurate at predicting specific values but are also effective in replicating the overall behavior and volatility characteristics of the market.

Upon closer examination, the GRU model's predicted distribution shows a slightly tighter alignment with the actual distribution curve compared to that of the LSTM model. The GRU's prediction density more accurately follows the peak and spread of the actual price distribution, suggesting that it captured the inherent volatility of the market with higher precision. This closer fit indicates that GRU was better able to generalize from historical price patterns and reproduce the statistical properties of unseen data. In contrast, the LSTM model, while still performing well, produced a distribution with marginally wider dispersion, which suggests that its predictions were slightly more variable around the mean price range. These differences further support the quantitative findings that GRU achieved lower RMSE and MAE values, reinforcing its effectiveness in modeling both the accuracy and variability of stock prices. Overall, the distribution analysis confirms that GRU provides a more reliable approximation of real market dynamics, making it a robust choice for financial forecasting applications.

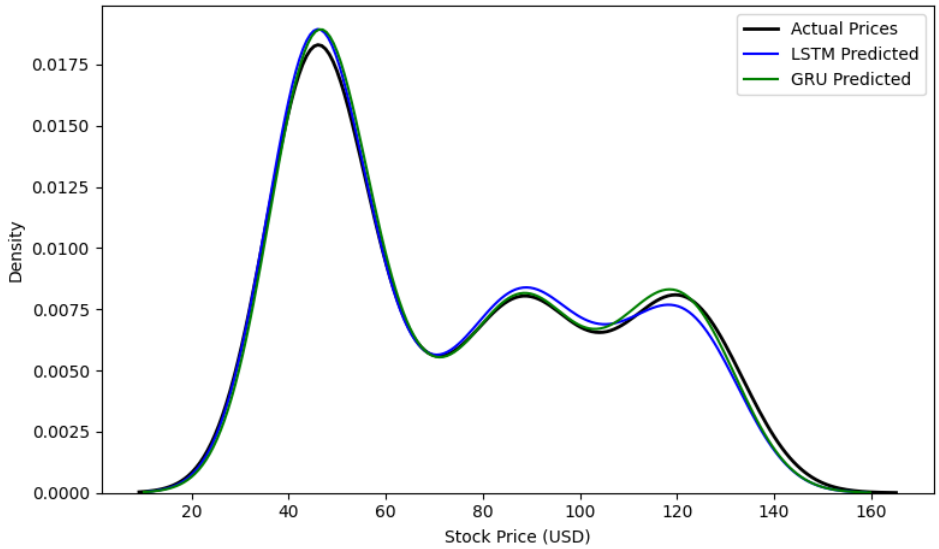


Figure 5 Distribution comparison of actual vs predicted NVIDIA stock prices.

Overall, these results confirm that both deep learning models are capable of effectively learning and predicting complex nonlinear patterns in financial time-series data. The GRU model, with its lower error metrics and faster convergence, is more suitable for real-time forecasting in dynamic digital financial markets, whereas the LSTM model demonstrates strong potential for applications requiring stable long-term trend prediction. The performance comparison between LSTM and GRU in this study reinforces the applicability of deep learning methods for improving accuracy and efficiency in AI-driven digital

market forecasting.

Discussion

The results of this study highlight the effectiveness of deep learning models, particularly Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), in predicting stock price movements within the context of the digital financial market [21], [22]. Both models demonstrated the ability to learn and capture complex temporal dependencies in NVIDIA's stock price data, producing accurate and stable forecasts [23]. The GRU model, however, exhibited superior performance in terms of predictive accuracy and training efficiency. Its lower RMSE and MAE values, coupled with faster convergence, indicate that GRU was better at adapting to short-term market fluctuations while maintaining prediction stability. This advantage can be attributed to GRU's simplified gating structure, which reduces computational complexity and facilitates efficient information flow during backpropagation [24]. The LSTM model also performed well, particularly in modeling long-term dependencies, although its more complex memory cell architecture may require additional computational resources and longer training times [25]. These findings are consistent with previous research on financial forecasting, which highlights the potential advantages of GRU in computational efficiency while recognizing the effectiveness of LSTM in modeling complex temporal patterns [26].

The comparative results also provide deeper insights into how these models behave under different market conditions. The GRU model's ability to capture abrupt price movements makes it a promising approach for environments characterized by high volatility, including short-term trading and digital asset forecasting [27]. Meanwhile, the LSTM model's smoother prediction curves may make it useful for identifying broader market trends that evolve over longer periods. Both models showed strong generalization capabilities, as indicated by residuals distributed symmetrically around zero and close alignment between predicted and actual price distributions. These observations suggest that the models effectively captured the underlying data patterns without substantial systematic prediction bias. In the context of digital finance, these findings underscore the potential of deep learning algorithms to enhance predictive analytics, improve automated trading systems, and support data-driven decision-making in dynamic and information-rich market environments [28].

Conclusion

This study concludes that both Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models are effective deep learning architectures for forecasting NVIDIA stock prices, with each model demonstrating distinct strengths in capturing the temporal dynamics of financial market behavior. The GRU model produced lower RMSE and MAE values, indicating higher predictive accuracy and faster convergence compared to LSTM, which suggests that GRU is more efficient in adapting to short-term fluctuations and computationally lighter for real-time applications. Meanwhile, the LSTM model maintained strong performance in modeling long-term dependencies, producing smoother prediction curves that are better suited for long-horizon trend analysis. Both models exhibited robust generalization capabilities, as seen in the close alignment between predicted and actual price distributions and the balanced residuals around zero. These findings underscore the potential of deep learning

techniques to improve forecasting accuracy in digital financial markets by effectively capturing nonlinear and sequential patterns in stock data. However, as this study relied primarily on historical price and volume data, future research should integrate additional market indicators, such as sentiment data, macroeconomic variables, or news-based information, to further enhance predictive robustness and provide more comprehensive insights into market dynamics.

Declarations

Author Contributions

Conceptualization: R.B.; Methodology: R.B.; Software: R.B.; Validation: R.B.; Formal Analysis: R.B.; Investigation: R.B.; Resources: R.B.; Data Curation: R.B.; Writing Original Draft Preparation: R.B.; Writing Review and Editing: R.B.; Visualization: R.B.; The author has read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] R. F. Engle, “Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation,” *Econometrica*, vol. 50, no. 4, pp. 987–1007, 1982, doi: 10.2307/1912773.

[2] T. Bollerslev, “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, vol. 31, no. 3, pp. 307–327, 1986, doi: 10.1016/0304-4076(86)90063-1.

[3] G. T. Wilson, “Review of Time Series Analysis: Forecasting and Control, 5th Edition, by George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung,” *Journal of Time Series Analysis*, vol. 37, no. 5, pp. 709-711, Mar. 2016, doi: 10.1111/jtsa.12194.

[4] J. Patel, S. Shah, P. Thakkar, and K. Kotecha, “Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques,” *Expert Systems with Applications*, vol. 42, no. 1, pp. 259–268, 2015,

doi: 10.1016/j.eswa.2014.07.040.

[5] W. Huang, Y. Nakamori, and S. Y. Wang, “Forecasting stock market movement direction with support vector machine,” *Computers & Operations Research*, vol. 32, no. 10, pp. 2513–2522, 2005, doi: 10.1016/j.cor.2004.03.016.

[6] W. Bao, J. Yue, and Y. Rao, “A deep learning framework for financial time series using stacked autoencoders and long short-term memory,” *PLOS ONE*, vol. 12, no. 7, p. e0180944, 2017, doi: 10.1371/journal.pone.0180944.

[7] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.

[8] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994, doi: 10.1109/72.279181.

[9] K. Cho et al., “Learning phrase representations using RNN encoder–decoder for statistical machine translation,” in *Proc. EMNLP*, Doha, Qatar, vol. 2014, no. October, pp. 1724–1734, doi: 10.3115/v1/D14-1179.

[10] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *arXiv preprint arXiv:1412.3555*, vol. 2014, no. December, pp. 1-9, 2014. Online. Available: https://arxiv.org/abs/1412.3555.

[11] T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” *European Journal of Operational Research*, vol. 270, no. 2, pp. 654–669, 2018, doi: 10.1016/j.ejor.2017.11.054.

[12] D. M. Q. Nelson, A. C. M. Pereira, and R. A. de Oliveira, “Stock market’s price movement prediction with LSTM neural networks,” in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Anchorage, AK, USA, vol. 2017, no. July, pp. 1419–1426, 2017, doi: 10.1109/IJCNN.2017.7966019.

[13] X. Huang, “Stock price prediction based on neural networks incorporating attention mechanism,” in *Proc. 3rd Int. Conf. Internet Finance Digit. Economy (ICIFDE)*, vol. 2023, no. October, pp. 57–66, 2023, doi: 10.2991/978-94-6463-270-5_57.

[14] A. Sagheer and M. Kotb, “Time series forecasting of petroleum production using deep LSTM recurrent networks,” *Neurocomputing*, vol. 323, no. January, pp. 203–213, 2019, doi: 10.1016/j.neucom.2018.09.082.

[15] S. Siami-Namini, N. Tavakoli, and A. Siami Namin, “A comparison of ARIMA and LSTM in forecasting time series,” in *Proc. IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, Boca Raton, FL, USA, vol. 2019, no. January, pp. 1394–1401, 2018, doi: 10.1109/ICMLA.2018.00227.

[16] H. Baek, “A CNN–LSTM stock prediction model based on genetic algorithm optimization,” *Asia-Pacific Financial Markets*, vol. 31, no. June, pp. 205–220, 2023, doi: 10.1007/s10690-023-09412-z.

[17] Y. Khodke and S. Deshpande, “Stock Price Prediction Using LSTM and GRU,” *International Journal for Research in Applied Science and Engineering Technology (IJRASET)*, vol. 11, no. 6, pp. 1163–1167, Jun. 2023, doi: 10.22214/ijraset.2023.53826.

[18] Y. Nejatbakhsh and M. Aliasgari, “Enhancing Stock Market Prediction With Hybrid

Deep Learning: Integrating LSTM, Transformer Attention, Federated Learning, and Sentiment Analysis," *IEEE Access*, vol. 12, no. December, pp. 125876–125889, 2024, doi: 10.1109/ACCESS.2025.3649668.

[19] J. Sun, "Stock Price Prediction Using LSTM with Attention-Based Multimodal Fusion Model," *Economics & Management Information*, vol. 4, no. 4, pp. 1-5, Aug. 2025, doi: 10.62836/emi.v4i4.506.

[20] J. W. Taylor and R. Buizza, "Using weather ensemble predictions in electricity demand forecasting," *International Journal of Forecasting*, vol. 19, no. 1, pp. 57–70, 2003, doi: 10.1016/S0169-2070(01)00123-6.

[21] A. Farhadi, A. Zamanifar, A. Alipour, A. Taheri, and M. Asadolahi, "A hybrid LSTM-GRU model for stock price prediction," *IEEE Access*, vol. 13, no. 1, pp. 117594–117618, 2025, doi: 10.1109/ACCESS.2025.3586558.

[22] E. J. Prasetyo and K. Hartomo, "Multi-industry stock forecasting using GRU-LSTM deep transfer learning method," *J. Infotel*, vol. 15, no. 2, pp. 1–10, May 2023, doi: 10.20895/infotel.v15i2.941.

[23] R. Thulasiram, C. J. Henry, and A. Thavaneswaran, "Encoder-decoder based LSTM and GRU architectures for stocks and cryptocurrency prediction," *J. Risk Financial Manage.*, vol. 17, no. 5, pp. 1–10, May 2024, doi: 10.3390/jrfm17050200.

[24] M. J. Hamayel and A. Y. Owda, "A novel cryptocurrency price prediction model using GRU, LSTM and bi-LSTM machine learning algorithms," *AI*, vol. 2, no. 4, pp. 477–496, Dec. 2021, doi: 10.3390/ai2040030.

[25] A. Dutta, S. Kumar, and M. Basu, "A gated recurrent unit approach to Bitcoin price prediction," *J. Risk Financial Manage.*, vol. 13, no. 2, pp. 1–12, Feb. 2020, doi: 10.3390/jrfm13020023.

[26] M. Kolambe and S. Arora, "Comparative analysis of LSTM variants for stock price forecasting on NSE India: GRU's dominance and enhancements," *Int. J. Inf. Technol. Comput. Sci.*, vol. 2024, no. 6, pp. 1–10, Dec. 2024, doi: 10.5815/ijitcs.2024.06.04.

[27] N. Tripathy, S. Hota, D. Mishra, P. Satapathy, and S. K. Nayak, "Empirical forecasting analysis of Bitcoin prices," *Int. J. Electr. Comput. Eng. Syst.*, vol. 15, no. 1, pp. 1–10, Jan. 2024, doi: 10.32985/ijeces.15.1.3.

[28] A. S. SGirsang and Stanley, "Hybrid LSTM and GRU for cryptocurrency price forecasting based on social network sentiment analysis using FinBERT," *IEEE Access*, vol. 11, no. 1, pp. 120530–120540, 2023, doi: 10.1109/ACCESS.2023.3324535.