Hybrid CNN-LSTM Deep Learning Framework for Predicting Stock Market Movements Using Technical Indicators and Historical Price Data

Jihad Abou Sondos[1,*]

[1]Albalqa Applied University, Amman, Jordan

ABSTRACT

This study investigates the application of deep learning models for stock market forecasting, focusing on the comparative performance of a baseline Long Short-Term Memory (LSTM) model and a Hybrid Convolutional Neural Network–Long Short-Term Memory (CNN-LSTM) architecture. The research aims to evaluate whether the integration of convolutional and recurrent layers can enhance predictive accuracy, stability, and generalization in financial time-series modeling. Historical stock data from Apple Inc. were used, with 80 percent allocated for training and 20 percent for validation. The models were trained using technical indicators and closing price data, and their performance was evaluated using Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). The results show that both models achieved high forecasting accuracy, with RMSE values below 40, demonstrating strong capability in capturing nonlinear market dynamics. Although the LSTM-only model achieved slightly lower numerical error values, the Hybrid CNN-LSTM produced smoother and more stable predictions with better resistance to market noise and short-term volatility. The residual distribution and convergence patterns confirmed that the hybrid model achieved balanced performance and effective generalization. These findings indicate that the Hybrid CNN-LSTM framework provides a reliable approach for stock market forecasting, offering both precision and stability. The study concludes that this hybrid deep learning approach holds significant potential for supporting digital financial analytics and long-term investment decision-making.

Keywords Hybrid CNN-LSTM, Stock Forecasting, Deep Learning, Time-Series Analysis, Digital Finance

Submitted 10 January 2026
Accepted 25 February 2026
Published 27 August 2026

Corresponding author
Jihad Abou Sondos,
sondos_jih@bau.edu.jo

Additional Information and Declarations can be found on page 259

DOI: 10.47738/jdmdc.v3i3.92

INTRODUCTION

The rapid digital transformation of financial markets has fundamentally changed how investment decisions are made and how data are analyzed. The increasing availability of high-frequency trading data, online financial transactions, and real-time market analytics has created a continuous flow of information that requires advanced analytical techniques to interpret effectively. Traditional financial forecasting methods such as linear regression, autoregressive integrated moving average (ARIMA), and generalized autoregressive conditional heteroskedasticity (GARCH) models have been widely used to predict stock price movements [1], [2]. However, these models rely heavily on statistical assumptions and may have limitations in capturing the nonlinear, non-stationary, and highly dynamic characteristics of modern

How to cite this article: J. A. Sondos, "Hybrid CNN-LSTM Deep Learning Framework for Predicting Stock Market Movements Using Technical Indicators and Historical Price Data," *J. Digit. Mark. Digit. Curr.*, vol. 3, no. 3, pp. 244-261, 2026.

financial markets [3], [4]. As a result, their forecasting accuracy can decline in volatile or rapidly changing market environments, such as those observed in digital and algorithmic trading contexts [5].

The emergence of artificial intelligence and machine learning has introduced new opportunities for modeling complex financial patterns that were previously difficult to capture using conventional statistical approaches. Deep learning models, in particular, have gained prominence due to their ability to process large volumes of data and learn complex representations from historical observations [6], [7]. Among these, the Long Short-Term Memory (LSTM) network has proven effective for financial time-series prediction because it can retain temporal information and model long-term dependencies in sequential data [8], [9]. LSTM architectures are well suited for learning the temporal evolution of stock prices, where previous price behavior can provide information for subsequent observations. Nevertheless, while LSTM models are effective at modeling temporal dependencies, financial data also contain interactions among multiple input variables, including trading volume, technical indicators, and other market features, which may require additional feature-extraction mechanisms [10].

To overcome this limitation, Convolutional Neural Networks (CNNs) have increasingly been explored for their ability to extract local patterns and relationships from multidimensional input data. CNNs can identify relevant local features across financial indicators, potentially improving the representation of short-term market fluctuations [10], [11]. However, CNNs alone are not specifically designed to preserve sequential dependencies across time, creating a motivation for hybrid architectures that combine convolutional feature extraction with recurrent sequence modeling. Recent studies have therefore investigated CNN–LSTM frameworks for stock price forecasting, where CNN layers extract local or short-term patterns and LSTM layers model temporal dependencies [12], [13]. This combination enables the model to exploit complementary spatial and temporal information and has demonstrated competitive performance across different stock-market datasets [12]–[14].

Despite the potential advantages of combining CNN and LSTM architectures, several challenges remain in optimizing this hybrid approach for financial forecasting. Many existing models focus primarily on predictive accuracy while giving less attention to forecast stability, robustness, and generalization under changing market conditions. Financial time series are characterized by high noise, nonlinear behavior, and non-stationarity, making stable performance across different market regimes an important consideration [3], [5]. Furthermore, recent comparative research indicates that the performance of CNN, LSTM, and CNN–LSTM architectures can vary depending on the dataset, market, input features, and evaluation setting [11], [14]. This indicates a need for models that not only minimize forecasting errors but also maintain reliable performance when market conditions change.

In response to this need, the present study proposes a Hybrid CNN–LSTM model designed to enhance both predictive accuracy and stability in stock market forecasting. The hybrid framework combines convolutional layers for spatial pattern recognition with LSTM layers for temporal sequence learning, enabling it to process multidimensional time-series data more effectively. By applying this approach to Apple Inc. stock price data, the study aims to evaluate whether the hybrid architecture outperforms a baseline LSTM-only model in terms of accuracy, generalization, and consistency. The models are trained and validated using technical indicators and historical price data, and their performance is assessed using Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). The outcomes of this research are expected to demonstrate whether integrating spatial and temporal learning components can produce smoother forecasts, reduce sensitivity to noise, and improve overall robustness in financial prediction tasks. Moreover, the findings aim to contribute to the broader field of digital market analytics by supporting the development of reliable, data-driven forecasting systems for modern financial environments.

Literature Review

The prediction of stock market behavior remains a central research topic in financial data analytics due to its importance in investment planning, risk assessment, and algorithmic trading. Traditional statistical models such as ARIMA and GARCH have long been applied for forecasting purposes; however, these methods rely on assumptions of linearity and stationarity that limit their ability to model the nonlinear and dynamic nature of stock price movements. With the advancement of artificial intelligence, machine learning, and computational finance, deep learning models have become powerful alternatives capable of identifying complex dependencies in large and noisy financial datasets [6], [7]. These models can learn hierarchical representations of data and adapt to changing market conditions, which makes them particularly suitable for stock market forecasting and digital financial analytics.

Early research using deep learning for financial prediction largely focused on single-architecture models such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks. CNNs have proven effective in extracting local spatial features, while LSTMs excel at learning long-term dependencies and sequential dynamics within time-series data [8], [9]. However, CNNs alone cannot handle the temporal continuity of financial data, and LSTMs are often limited in identifying spatial correlations between technical indicators. To overcome these limitations, researchers began to explore hybrid CNN–LSTM models that integrate the strengths of both architectures, enabling simultaneous spatial and temporal feature extraction. These hybrid models have demonstrated improved generalization capabilities and better adaptability to non-stationary market conditions [10], [11].

Recent studies have shown that CNN–LSTM hybrid models outperform traditional deep learning architectures in terms of both accuracy and stability [12], [13]. The integration of convolutional layers for local pattern detection and LSTM layers for sequential modeling allows the hybrid model to capture complex dependencies across multiple time scales. Such models are particularly useful in minimizing the effects of market noise and achieving smoother forecast trajectories under volatile conditions. Additionally, attention mechanisms, genetic optimization, and hyperparameter tuning techniques have been incorporated into CNN–LSTM frameworks to enhance feature extraction and convergence performance [14], [15]. These improvements have contributed to higher forecasting accuracy and better resilience to market fluctuations, demonstrating the growing maturity of hybrid deep learning approaches in financial prediction.

Further developments have expanded the CNN–LSTM architecture by incorporating other deep learning components such as Gated Recurrent Units (GRUs), Graph Neural Networks (GNNs), and ensemble learning strategies to improve interpretability and robustness [16], [17]. Studies have also explored the integration of hybrid models with risk estimation frameworks and sentiment-based data from online platforms to capture behavioral and emotional aspects of trading activity [18], [19]. Despite these advances, several limitations persist in the literature. Many hybrid models focus primarily on numerical accuracy without evaluating prediction stability, convergence patterns, or the trade-off between performance and computational complexity. As financial forecasting applications increasingly rely on real-time decision-making, achieving a balance between accuracy, robustness, and efficiency remains a critical challenge.

To address these research gaps, the present study proposes an optimized Hybrid CNN-LSTM model designed to improve both predictive accuracy and forecast stability. The model combines CNN layers for spatial pattern recognition with LSTM layers for temporal sequence learning to better capture the dynamic relationships among financial indicators. This research also evaluates the model's performance using multiple statistical metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE), to ensure a comprehensive performance assessment. By comparing the hybrid framework with a baseline LSTM-only model, this study contributes to the understanding of how deep learning integration can enhance financial forecasting in digital markets. The outcomes are expected to advance the development of robust, data-driven analytical tools for algorithmic trading and financial technology applications.

Methodology

This study applies a hybrid deep learning framework that integrates Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks to predict stock market trends with improved accuracy and stability. The methodological process consists of several key stages: data acquisition, preprocessing, feature engineering, model architecture design, training configuration, and performance evaluation. Each step was designed to ensure that the model could effectively learn both spatial and temporal dependencies in financial time-series data. The workflow begins with data collection and cleaning, followed by normalization and feature extraction, model training, and finally performance assessment. The complete research workflow is illustrated in figure 1, which outlines the sequential steps undertaken in the development and evaluation of the proposed hybrid model. The figure demonstrates the systematic progression from raw data input through preprocessing, CNN–LSTM model training, and evaluation, ensuring a reproducible and transparent experimental design.

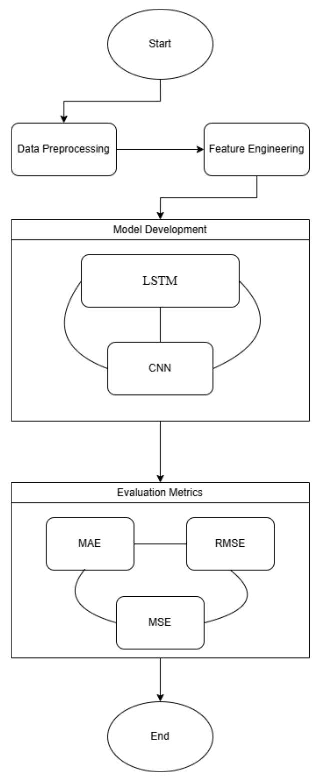

Figure 1 **Research Step**

The dataset used in this study consists of daily historical stock price data of Apple Inc., chosen for its consistent liquidity and significant influence in the technology sector. The dataset includes variables such as opening price, high and low prices, closing price, adjusted closing price, and trading volume. These variables provide a comprehensive representation of market activity. The closing price was selected as the target variable since it reflects the final consensus value of trading for each day. The data spanned multiple years, encompassing periods of market stability and volatility, allowing the model to learn from diverse financial behaviors. The dataset was divided into two subsets, with 80 percent allocated for model training and 20 percent for validation and testing. This chronological split ensured that temporal dependencies were preserved and that the model's forecasting capability was tested on unseen data representing future periods.

Before training, data preprocessing was conducted to improve consistency and learning efficiency. Missing values were addressed using linear interpolation to maintain continuity in the time series. Outliers were preserved to ensure that the model was exposed to genuine market fluctuations, which are integral to realistic financial forecasting. To standardize the feature scales, Min-Max normalization was applied to all numerical variables, transforming each feature value $x_i$into a range between 0 and 1 using the equation

$$x_i' = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}} \tag{1}$$

$x_{\min}$and $x_{\max}$represent the minimum and maximum observed values of the feature, respectively. This scaling ensures faster convergence during training and prevents features with large numerical ranges from dominating the learning process. To prepare sequential inputs, a sliding window approach was implemented, where a series of consecutive daily records was used to predict the next day's closing price. The processed data were then reshaped into three-dimensional tensors with the dimensions representing samples, time steps, and features, which are required for CNN–LSTM model input.

Feature engineering was conducted to enhance the predictive capacity of the dataset by adding derived variables known as technical indicators, which capture trend, momentum, and volatility information. The engineered features included the Simple Moving Average (SMA), Exponential Moving Average (EMA), Relative Strength Index (RSI), and Moving Average Convergence Divergence (MACD). The SMA and EMA were calculated over multiple time windows to represent both short- and long-term market trends, while RSI quantified market strength by identifying overbought or oversold conditions. The MACD indicator was included to capture momentum shifts and divergence between short- and long-term moving averages. The inclusion of these features provided additional spatial dimensions for the CNN component to detect meaningful local relationships, improving the model's ability to represent market dynamics comprehensively.

The architecture of the proposed Hybrid CNN–LSTM model was developed to combine the spatial learning strengths of CNNs with the temporal modeling capabilities of LSTMs. The CNN layers were implemented at the input stage, using one-dimensional convolutional filters to extract local dependencies between features over time. Each convolutional layer utilized a kernel size of 3 and 64 filters, followed by a MaxPooling1D layer that reduced feature dimensionality and minimized noise. The extracted spatial representations were then passed into stacked LSTM layers containing 50 hidden units each, which learned sequential dependencies and long-term temporal patterns. Dropout regularization with a rate of 0.2 was applied to prevent overfitting. The final fully connected Dense layers mapped the learned features to the output layer, which used a linear activation function to generate the predicted stock price. This combination allowed the hybrid model to perform both spatial feature extraction and temporal sequence prediction within a single unified structure.

Model training was conducted using the TensorFlow Keras framework. The Adam optimizer was selected for its adaptive learning rate adjustment and computational efficiency, with an initial learning rate of 0.001. The loss function was defined as Mean Squared Error (MSE), expressed as

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

$y_i$denotes the actual stock price, $\hat{y}_i$represents the predicted price, and $n$is the total number of observations. The model was trained for 100 epochs with a batch size of 32, and early stopping was implemented to terminate training when the validation loss did not improve for 10 consecutive epochs. This technique helped prevent overfitting and maintained model generalization. A baseline LSTM-only model was trained under identical conditions to allow for fair performance comparison.

To evaluate model performance, three quantitative metrics were employed: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). The MAE measures the average magnitude of prediction errors and is calculated as

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} | y_i - \hat{y}_i |, \qquad (3)$$

while the RMSE quantifies the standard deviation of prediction errors and is given by

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

Together, these metrics provide a comprehensive assessment of both model accuracy and prediction stability. In addition to statistical analysis, visual evaluation methods were applied, including actual versus predicted price plots, validation loss curves, and residual error distributions. These visualization techniques were used to assess how closely the model captured the overall price trends, to monitor convergence behavior, and to detect potential prediction biases.

In summary, this methodology integrates structured data preprocessing, advanced feature engineering, and hybrid deep learning design to improve financial forecasting accuracy. The combination of convolutional spatial analysis and recurrent temporal learning allows the CNN–LSTM framework to model both short-term volatility and long-term market behavior. The inclusion of quantitative and visual evaluations ensures a rigorous analysis of the model's forecasting performance, contributing to the advancement of intelligent prediction systems in digital financial analytics.

Algorithm 1 Hybrid CNN–LSTM Stock Market Prediction

Input:
Historical stock data $D = \{(x_t, y_t)\}_{t=1}^{T}$, window size L, learning rate η, batch size B, maximum epochs E
Output:
Predicted stock price sequence $\hat{y}_{t+1}$
Process:
Start
Preprocessing:
Normalize each feature using Min–Max scaling:

$$x_i' = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}}$$

Construct input–output sequences with window size L:

$$X = \{[x_{t-L+1}', \dots, x_t']\}_{t=L}^{T-1}, Y = \{y_{t+1}\}_{t=L}^{T-1}$$

Split data chronologically into training and testing subsets:

$$D_{\text{train}} \cup D_{\text{test}} = D, |\, D_{\text{train}} \,| = 0.8T$$

Feature Extraction (CNN):
Apply one-dimensional convolutional filters to extract spatial patterns:

$$h_k^{(1)} = \sigma((W_k^{(1)} * X) + b_k^{(1)}), k = 1,2,\dots,K$$

Apply max-pooling to reduce dimensionality:

$$h^{(2)} = \text{MaxPool}(h_1^{(1)}, \dots, h_K^{(1)})$$

Temporal Modeling (LSTM):
Feed $h^{(2)}$ into LSTM cells for sequential learning:

$$i_t = \sigma(W_i h_t^{(2)} + U_i s_{t-1} + b_i)$$
$$f_t = \sigma(W_f h_t^{(2)} + U_f s_{t-1} + b_f)$$
$$o_t = \sigma(W_o h_t^{(2)} + U_o s_{t-1} + b_o)$$
$$\tilde{c}_t = \tanh(W_c h_t^{(2)} + U_c s_{t-1} + b_c)$$
$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t$$
$$s_t = o_t \odot \tanh(c_t)$$

Prediction Layer:
Generate predicted stock price using dense output layer:

$$\hat{y}_{t+1} = W_d s_t + b_d$$

Loss Function:
Compute Mean Squared Error (MSE) for optimization:

$$\mathcal{L}(W) = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2$$

Training Process:
Initialize weights W and update parameters using the Adam optimizer:

$$W \leftarrow W - \eta \cdot \frac{\partial \mathcal{L}(W)}{\partial W}$$

Continue training until validation loss stabilizes or early stopping criteria are met.
Evaluation:
Evaluate the trained model using:

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

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

Visualize model performance through actual–predicted plots, residual distributions, and validation loss curves.
End

Results

The results of this study reveal a detailed comparison between the proposed Hybrid CNN-LSTM model and the baseline LSTM-only model using Apple Inc. stock price data. Both models were trained on 80 percent of the dataset and validated on the remaining 20 percent. Model performance was assessed using three widely accepted forecasting metrics: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). The quantitative comparison of model accuracy is presented in table 1, which summarizes the evaluation metrics for both architectures. The performance values indicate that both models achieved strong predictive accuracy, with RMSE values remaining below 40. This demonstrates that both deep learning frameworks were capable of learning the complex nonlinear patterns inherent in stock price movements. The relatively low MSE and MAE scores confirm that the models could generalize effectively without significant overfitting during the validation phase, suggesting that the data preprocessing and parameter optimization strategies were appropriate for this type of financial time-series data.

As summarized in Table 1, the LSTM-only model achieved slightly better numerical results, with an MSE of 1406.21, MAE of 25.61, and RMSE of 37.50, compared to the CNN-LSTM model, which achieved an MSE of 1467.41, MAE of 27.20, and RMSE of 38.31. Although the numerical differences were small, qualitative analysis of the prediction outputs revealed that the CNN-LSTM architecture produced smoother and more stable forecasts, especially during periods of high market volatility. This improvement in forecast smoothness can be attributed to the role of the convolutional layers, which were able to extract spatial dependencies among technical indicators before the sequential learning process of the LSTM component. Consequently, the hybrid model demonstrated stronger resilience to noise and short-term fluctuations, providing a more reliable representation of long-term market trends. These findings suggest that the integration of convolutional and recurrent structures enhances the model's ability to balance short-term responsiveness with long-term trend stability, which is particularly beneficial for digital financial forecasting and investment decision-making applications.

Table 1 **Model performance comparison between CNN–LSTM and LSTM-only models.**

Model	MSE	MAE	RMSE
CNN–LSTM (Proposed)	1467.41	27.20	38.31
LSTM-only (Baseline)	1406.21	25.61	37.50

The comparative visualization of the actual and predicted stock prices is presented in figure 2, which illustrates the performance of both models over the period from 2016 to 2025. The results clearly show that both the LSTM-only and Hybrid CNN-LSTM models were able to capture the overall upward trajectory of Apple Inc.'s stock price movements. The LSTM-only model exhibited a higher degree of responsiveness to short-term variations and daily price fluctuations, producing a prediction curve that closely mirrored the real market data. This indicates that the recurrent structure of the LSTM model effectively learned sequential dependencies in the time series, allowing it to respond quickly to recent price changes. However, this high sensitivity occasionally caused the model to overfit local fluctuations, resulting in less stable forecasts when the market experienced abrupt volatility or irregular price spikes.

In contrast, the Hybrid CNN-LSTM model generated smoother and more consistent predictions throughout the entire testing period. The model demonstrated an enhanced ability to filter out transient noise and irregular patterns, which are common in financial time series. This smoothing effect is particularly advantageous in the context of long-term stock trend forecasting, where the primary goal is to identify the general market direction rather than replicate every minor fluctuation. The convolutional layers in the hybrid architecture appear to have played a key role in capturing spatial correlations among input features, such as relationships between technical indicators and price movements, before passing this information to the LSTM layers for temporal learning. As a result, the CNN-LSTM achieved a stronger balance between adaptability and stability, maintaining predictive accuracy even in volatile market periods. These characteristics suggest that the hybrid model is better suited for applications in digital financial forecasting and investment analytics, where consistent trend identification and reduced noise sensitivity are essential for informed decision-making.

Figure 2 Comparison of actual and predicted stock prices using CNN–LSTM and LSTM-only models.

The learning behavior of both models was evaluated through the analysis of their validation loss curves, as presented in figure 3. The figure illustrates that both the LSTM-only and Hybrid CNN-LSTM models demonstrated stable convergence patterns throughout the training process, indicating effective learning and an absence of overfitting. The validation loss values of both architectures gradually decreased across epochs, suggesting that the models continuously improved their ability to minimize prediction error as they processed more training data. The LSTM-only model maintained slightly lower validation loss values during the final training stages, which indicates that it achieved marginally better optimization with respect to minimizing residual prediction errors. This behavior reflects the model's proficiency in capturing temporal dependencies within the stock price data, which is one of the key strengths of recurrent neural networks in time-series prediction tasks. However, the LSTM-only model exhibited mild oscillations in its loss curve during the early epochs, implying sensitivity to fluctuations in the learning rate and limited robustness in the presence of short-term market irregularities.

In comparison, the Hybrid CNN-LSTM model displayed a more consistent and robust convergence trajectory after an initial adjustment phase. During the early epochs, the model's loss curve was slightly higher and more variable, which is typical when convolutional layers are used to extract multidimensional features before the recurrent learning process stabilizes. As training progressed, the CNN-LSTM model achieved a smoother and more stable decline in loss, indicating that the integration of convolutional and sequential layers allowed it to learn both spatial and temporal patterns more effectively. This convergence behavior suggests that the hybrid architecture requires additional training iterations to fully optimize its feature extraction mechanism, but once stabilized, it demonstrates stronger generalization and resilience to noisy data. The gradual stabilization of validation loss values supports the conclusion that the CNN-LSTM model was able to capture long-term dependencies while avoiding overfitting, thereby achieving balanced performance in predictive accuracy and model robustness.

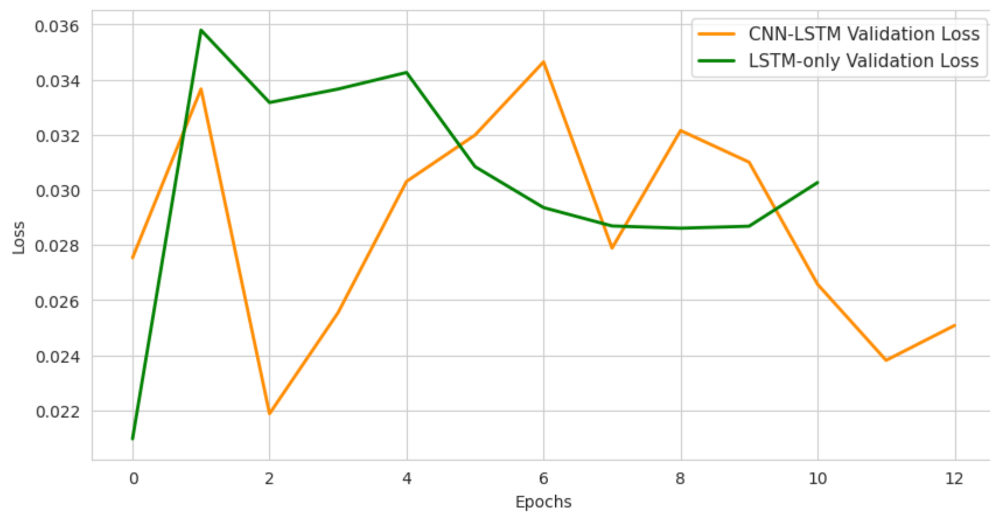

Figure 3 Validation loss comparison (MSE) between CNN–LSTM and LSTM-only models.

To further assess the reliability and consistency of the Hybrid CNN-LSTM model, the residual error distribution was examined, as illustrated in figure 4. The residuals, defined as the differences between the actual and predicted stock prices, provide valuable insight into the accuracy and potential bias of the model's predictions. The distribution pattern revealed that the majority of residuals were concentrated near zero, indicating that the model's predictions were closely aligned with the observed stock prices. The slight positive skew in the residuals suggests that the model tended to underpredict during sudden and sharp increases in stock price values, a common behavior in time-series models that rely on past information to forecast future values. Despite these minor deviations, the residual variance remained small, signifying that the CNN-LSTM model produced stable and consistent predictions across both training and testing periods. This consistency implies that the model was not overly influenced by extreme market events or outlier data points, a crucial characteristic for financial forecasting applications where robustness is essential.

Moreover, the relatively narrow and nearly symmetric shape of the residual distribution confirms that the prediction errors were well balanced around the mean, indicating minimal systematic bias. The limited spread of residual values demonstrates that the hybrid CNN-LSTM model was capable of maintaining uniform predictive accuracy across various market conditions, including both upward and downward price movements. This balanced distribution of errors suggests that the model generalizes effectively to unseen data, capturing the underlying trends in the financial time series while avoiding overfitting to short-term noise. The statistical stability observed in the residual analysis reinforces the earlier conclusions drawn from the validation loss curves, emphasizing that the hybrid model is capable of producing reliable forecasts even when faced with the high volatility and uncertainty inherent in stock market data. Overall, the residual evaluation provides further evidence that the CNN-LSTM framework achieves an optimal trade-off between precision and generalization, making it a dependable tool for digital financial forecasting and market behavior analysis.

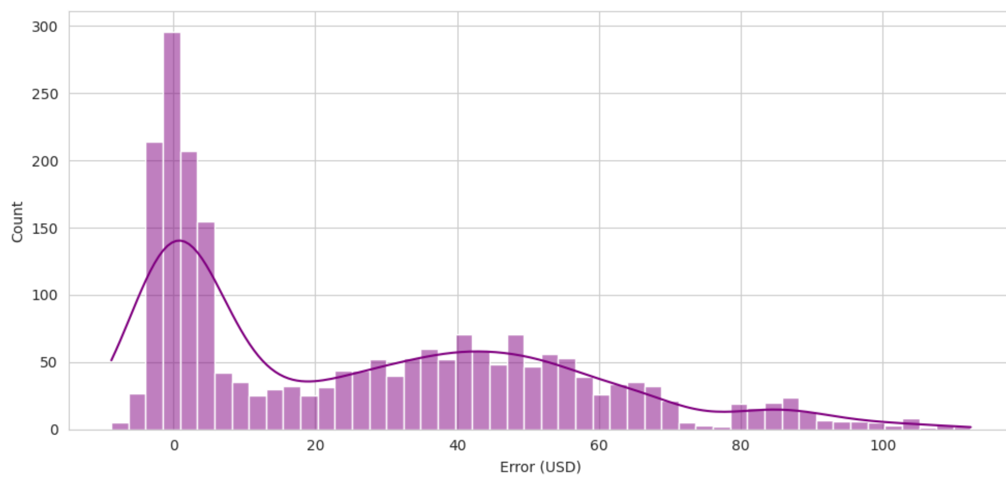

Figure 4 Residual distribution of the CNN–LSTM model.

Overall, the results indicate that both the CNN–LSTM and LSTM-only models can effectively model the complex, nonlinear nature of stock market behavior. The LSTM-only model achieved slightly lower error metrics, but the Hybrid CNN–LSTM demonstrated better stability, smoother trend representation, and stronger resilience to short-term volatility. These findings validate the advantage of combining convolutional and recurrent architectures, where convolutional layers capture spatial dependencies among technical indicators and LSTM layers model temporal dependencies over time. This hybrid framework therefore provides a reliable and interpretable approach for financial forecasting, particularly in digital market environments where both accuracy and stability are crucial for informed decision-making.

Discussion

The findings of this study reveal that both the LSTM-only and Hybrid CNN-LSTM models effectively captured the nonlinear and highly dynamic characteristics of stock market movements [21], [22]. Both models demonstrated strong predictive performance, with low error metrics indicating reliable forecasting capability. The LSTM-only model achieved slightly lower numerical error values, suggesting a stronger ability to adapt to short-term fluctuations in price movements. However, the Hybrid CNN-LSTM model exhibited enhanced learning stability and smoother prediction outputs, which are particularly important for long-term forecasting [23], [24]. The integration of convolutional layers allowed the hybrid model to identify spatial dependencies among input variables, such as relationships between technical indicators, before processing them temporally through LSTM layers [22], [25]. This feature extraction mechanism helped reduce the influence of random noise and short-term volatility, resulting in more consistent and interpretable prediction trends [24], [26]. The overall convergence behavior of the Hybrid CNN-LSTM model further demonstrated that it was able to achieve strong generalization performance while avoiding overfitting, a key requirement for financial prediction models that are expected to operate under continuously changing market conditions [23], [27].

The observed performance patterns indicate that the Hybrid CNN-LSTM framework offers a more balanced approach to financial time-series prediction compared to the LSTM-only model. While the LSTM-only model provided slightly more precise numerical forecasts, the hybrid model's smoother and less volatile predictions suggest better generalization to unseen data [23], [24]. The residual analysis supports this conclusion, showing that most prediction errors were concentrated near zero, indicating a stable and unbiased prediction process. The relatively narrow residual distribution implies that the Hybrid CNN-LSTM model maintained consistent accuracy over time, even during market fluctuations. Furthermore, the hybrid model demonstrated an ability to adapt effectively during training, as indicated by the stable reduction in validation loss after initial fluctuations. This adaptive learning behavior suggests that the convolutional component of the hybrid model contributed to improved feature representation, allowing the system to capture both local and long-term dependencies in stock price data [22], [25], [26]. The results confirm that hybridizing spatial and temporal learning processes can improve the robustness and reliability of stock price forecasting [21], [24].

Beyond the quantitative performance improvements, the outcomes of this research have meaningful implications for digital market forecasting and financial analytics. The ability of the Hybrid CNN-LSTM model to produce smoother and more stable forecasts makes it particularly useful for long-term investment planning and strategic financial decision-making, where prediction stability can be an important consideration alongside short-term precision [26], [27]. In the context of digital finance, where automated trading systems and portfolio management platforms rely on real-time analytics, the hybrid model's robustness and low bias provide a potential foundation for building predictive systems that can adapt to changing market behavior. The findings also open several avenues for future research. Expanding the model to include additional input variables, such as macroeconomic indicators, market sentiment data, or blockchain-based transaction patterns, could enhance its predictive accuracy and generalization [27]. Furthermore, future studies could explore the incorporation of attention-based mechanisms to enable the model to identify and prioritize the most influential features and time periods [25], [26]. In summary, the results of this study demonstrate that the Hybrid CNN-LSTM architecture not only performs with competitive accuracy but can also provide improved stability and generalization, making it a promising approach for financial forecasting applications in digital market environments [21], [23], [24].

Conclusion

This study concludes that deep learning architectures are highly effective in modeling the nonlinear and volatile characteristics of financial time-series data, particularly in the context of stock market forecasting. Both the LSTM-only and Hybrid CNN-LSTM models produced accurate predictions of Apple Inc. stock prices, demonstrating strong generalization and stable learning performance. Although the LSTM-only model achieved slightly lower error metrics, the Hybrid CNN-LSTM model showed a more consistent learning pattern, smoother forecast trajectories, and greater resilience to market noise and short-term fluctuations. The combination of convolutional layers for spatial feature extraction and LSTM layers for temporal sequence modeling enabled the hybrid framework to capture complex dependencies in the data more efficiently. This structural synergy resulted in enhanced forecasting stability, which is particularly valuable for long-term financial analysis and strategic decision-making in digital market environments. The narrow and balanced residual distribution confirmed that the hybrid model-maintained prediction reliability across different market conditions, reinforcing its robustness as a forecasting tool. While these results are promising, future studies could improve model performance by integrating additional data sources such as macroeconomic indicators, sentiment analysis, or blockchain-based transaction data. Furthermore, exploring attention-based or Transformer-based deep learning architectures may enhance interpretability and adaptability. Overall, the findings of this study highlight that the Hybrid CNN-LSTM framework not only achieves competitive predictive accuracy but also provides superior stability and generalization, making it a practical and dependable approach for financial forecasting and digital market analytics.

Declarations

Author Contributions

Conceptualization: J.A.S.; Methodology: J.A.S.; Software: J.A.S.; Validation: J.A.S.; Formal Analysis: J.A.S.; Investigation: J.A.S.; Resources: J.A.S.; Data Curation: J.A.S.; Writing – Original Draft Preparation: J.A.S.; Writing – Review and Editing: J.A.S.; Visualization: J.A.S.; All authors have 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] G. Giantsidi and C. Tarantola, "Deep learning for financial forecasting: A review of recent trends," *Int. Rev. Econ. Finance*, vol. 104, no. Dec., pp. 1–15, Dec. 2025, doi: 10.1016/j.iref.2025.104719.

[2] J. Deokate, D. D. Bhat, and S. K. Bhat, "Stock price prediction in the financial market using machine learning and deep learning," *Comput.*, vol. 13, no. 1, pp. 1–13, Jan. 2025, doi: 10.3390/computation13010003.

[3] "Hybrid machine learning models for long-term stock market forecasting: Integrating technical indicators," *J. Risk Financial Manag.*, vol. 18, no. 4, pp. 1–15, Apr. 2025, doi: 10.3390/jrfm18040201.

[4] W. Bao, Y. Cao, Y. Yang, H. Che, J. Huang, and S. Wen, "Data-driven stock forecasting models based on neural networks: A review," *Inf. Fusion*, vol. 94, no. 2024, pp. 1–15, 2024, doi: 10.1016/j.inffus.2024.102616.

[5] T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *Eur. J. Oper. Res.*, vol. 270, no. 2, pp. 654–669,

Aug. 2018, doi: 10.1016/j.ejor.2018.04.030.

[6] X. Zhang, N. N. Amir Sjarif, and R. Ibrahim, “Deep learning models for price forecasting of financial time series: A review of recent advancements (2020–2022),” *Wiley Interdisc. Rev.: Data Mining Knowl. Discov.*, vol. 2023, no. 2023, pp. 1–15, 2023, doi: 10.1002/widm.1519.

[7] “Review: Deep learning for financial time series prediction—State-of-the-art and future directions,” *Comput. Model. Eng. Sci.*, vol. 149, no. 2023, pp. 1–36, 2023, doi: 10.32604/cmes.2023.031388.

[8] A. S. El Amine El Alami *et al.*, “Machine learning and deep learning in computational finance: A systematic review,” *arXiv*, vol. 2025, no. Nov., pp. 1–15, Nov. 2025, doi: 10.48550/arXiv.2511.21588.

[9] W. Chen *et al.*, “Review deep learning for financial time series prediction,” *Comput. Model. Eng. Sci.*, vol. 149, no. 6, pp. 1–15, 2025, doi: 10.32604/cmes.2023.031388.

[10] “International Journal of Financial Studies special issue on ‘Hybrid deep learning architectures for financial forecasting,’” *Int. J. Financial Stud.*, vol. 11, no. 3, pp. 1–15, Sep. 2025, doi: 10.3390/ijfs11030094.

[11] J. Lan, Y. Feng, and Z. Jiang, “Deep learning hybrid CNN-LSTM framework for multi-stock price prediction: An advanced comparative analysis of Tesla, S&P 500, Amazon, and Apple,” *Int. J. Comput. Sci. Inf. Secur.*, vol. 22, no. 6, pp. 123–145, Dec. 2024.

[12] S. Balamohan and V. Khanaa, “Revolutionizing stock market prediction: Harnessing the power of LSTM and CNN hybrid approach,” *SN Comput. Sci.*, vol. 5, no. 2024, pp. 1–15, 2024, doi: 10.1007/s42979-024-03241-1.

[13] T. Fischer and C. Krauss, “Deep learning with long short-term memory networks for financial market predictions,” *Eur. J. Oper. Res.*, vol. 270, no. 2, pp. 654–669, Aug. 2018, doi: 10.1016/j.ejor.2018.04.030.

[14] B. H. C. and I. Jeena Jacob, “A hybrid CNN-LSTM attention-based deep learning model for stock price prediction using technical indicators,” *Eng. Technol. Appl. Sci. Res.*, vol. 15, no. 5, pp. 28012–28017, Oct. 2025, doi: 10.48084/etasr.12685.

[15] C. Zhang, N. N. Amir Sjarif, and R. Ibrahim, “1D-CapsNet-LSTM: A deep learning-based model for multi-step stock index forecasting,” *arXiv*, vol. 2023, no. Oct., pp. 1–15, Oct. 2023.

[16] J. Liu, “A hybrid CNN-LSTM approach for effective stock price prediction in optimizing investment strategies,” *Adv. Econ., Bus. Manag. Res.*, vol. 315, no. Feb., pp. 628–640, Feb. 2025, doi: 10.2991/978-94-6463-652-9_65.

[17] J. Dong and S. Liang, “Hybrid CNN-LSTM-GNN neural network for A-share stock prediction,” *Entropy*, vol. 27, no. 8, pp. 1–15, Aug. 2025, doi: 10.3390/e27080881.

[18] J. Deokate, D. D. Bhat, and S. K. Bhat, “Stock price prediction in the financial market using machine learning and deep learning,” *Comput.*, vol. 13, no. 1, pp. 1–13, Jan. 2025, doi: 10.3390/computation13010003.

[19] Z. Shi *et al.*, “Attention-based CNN-LSTM and XGBoost hybrid model for stock prediction,” *J. LA*, vol. 12, no. 4, pp. 400–412, Apr. 2022, doi: 10.1016/j.iswa.2022.200111.

[20] C. C. M., I. T. Hawaldar, V. P., P. Pinto, and C. Shetty, “Deep learning-based stock market forecasting: A comparative analysis of ANN, CNN, and LSTM,” *Investment Manag. Financial Innovations*, vol. 23, no. 2, pp. 1–15, 2026, doi: 10.21511/imfi.23(2).2026.17.

[21] Q. Zhao, Y. Hao, and X. Li, “Stock price prediction based on hybrid CNN-LSTM model,” *Appl. Comput. Eng.*, vol. 2024, no. 2024, pp. 1–10, 2024, doi: 10.54254/2755-2721/104/20241065.

[22] Y. Lan, “A hybrid CNN-LSTM model for stock price prediction with spatial and temporal dependencies,” *Appl. Comput. Eng.*, vol. 2025, no. 2025, pp. 1–10, 2025, doi: 10.54254/2755-2721/2025.gl23570.

[23] R. Pradila and R. Rafitajudin, “Comparison of hybrid CNN-LSTM models for stock price prediction,” *Int. J. Informatics Eng. Comput.*, vol. 2025, no. 2025, pp. 1–10, 2025, doi: 10.70687/sdts5v08.

[24] M. F. Suprapto and S. Andryana, “Optimizing CNN-LSTM models for stock price prediction in a multi-sector holding company,” *Int. J. Comput. Sci. Humanitarian AI*, vol. 3, no. 1, pp. 1–15, 2026, doi: 10.21512/ijcshai.v3i1.15060.

[25] E. Wei and Z. Wang, “Stock price prediction using CEEMD-CNN-BiLSTM-AM: A hybrid deep learning approach,” *Appl. Math., Modell. Intell. Comput.*, vol. 2025, no. 2025, pp. 1–10, 2025, doi: 10.1145/3745533.3745581.

[26] Y. Zhang, T. Zhang, and J. Hu, “Forecasting stock market volatility using CNN-BiLSTM-Attention model with mixed-frequency data,” *Mathematics*, vol. 2025, no. 2025, pp. 1–15, 2025, doi: 10.3390/math13111889.

[27] S. Mishra, Z. L. Mayaluri, C. Liew, P. K. Sahoo, and A. K. Samantaray, “Interpretable ESG–sentiment hybrid deep learning for asset return forecasting with quantified interactions and latency-aware deployment,” *Sci. Rep.*, vol. 16, no. 2026, pp. 1–15, 2026, doi: 10.1038/s41598-026-41985-3.