European Journal of AI, Computing & Informatics

Vol. 1 No. 3 2025

Article Open Access

Stock Price Prediction of Apple Inc. Based on LSTM Model: An Application of Artificial Intelligence in Individual Stock Analysis

Ruoyu Qi [1,*] and Liao Hu [2]

Received: 27 July 2025
Revised: 05 August 2025
Accepted: 16 August 2025
Published: 25 August 2025

[1] North Carolina State University, Raleigh, North Carolina, United States
[2] University of Illinois at Chicago, Chicago, United States
[*] Correspondence: Ruoyu Qi, North Carolina State University, Raleigh, North Carolina, United States

Abstract: This paper explores the potential of artificial intelligence in financial forecasting by applying a Long Short-Term Memory (LSTM) neural network to the task of predicting Apple Inc. (AAPL) stock prices. Using historical daily closing prices from August 2020 to August 2023, the model was trained and tested under a structured framework of data preprocessing, normalization, and sequential input construction. The forecasting performance was evaluated with widely used error measures, including Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). Experimental results indicate that the LSTM model effectively captures sequential dependencies and nonlinear dynamics within the time series, generating predictions that closely align with observed prices. While the model demonstrates high accuracy in relatively stable market conditions, its reliance on univariate input limits adaptability to abrupt market fluctuations and external influences such as macroeconomic shifts. These findings suggest that LSTM-based models can serve as valuable tools for supporting individual stock analysis, but their effectiveness depends on stock-specific data availability and market characteristics. Future work may extend this framework by incorporating multi-source financial indicators and enhancing model interpretability to achieve broader practical relevance.

Keywords: artificial intelligence; LSTM; stock price prediction; Apple Inc.; individual stock analysis; deep learning

1. Introduction

In recent years, the rising volatility and growing complexity of financial markets have drawn significant attention from both academia and industry toward stock price prediction. Achieving precise forecasts of stock movements has long been recognized as a difficult challenge in financial analysis, largely due to the stochastic and nonlinear characteristics of market dynamics. Conventional statistical techniques, including autoregressive models and moving averages, often fail to effectively capture the long-range dependencies and evolving patterns embedded in stock price fluctuations.

With the rapid progress of artificial intelligence (AI), particularly deep learning, scholars have started to investigate advanced algorithms capable of extracting temporal relationships and nonlinear behaviors from large-scale financial datasets. Among these approaches, Long Short-Term Memory (LSTM) networks—a specialized form of recurrent neural network (RNN)—have produced encouraging outcomes in time-series prediction by preserving and leveraging information from extended sequences. This property makes them especially well-suited for stock forecasting, where historical records frequently carry crucial signals about upcoming trends.

The integration of AI into financial modeling is not merely a technical innovation but represents a broader shift toward data-driven decision-making in the financial industry. Intelligent models, when properly trained and validated, can offer insights that support portfolio management, trading strategies, and risk assessment. The application of LSTM-based models to individual stock prediction exemplifies the practical convergence of machine learning and financial analysis.

This study focuses on Apple Inc. (AAPL) as the subject of analysis. Apple is one of the most traded and closely watched technology stocks in the global market, making it a representative case for evaluating AI-driven prediction methods. Its relatively high market liquidity, frequent news coverage, and strong correlation with market sentiment provide an ideal setting for testing the predictive capabilities of deep learning models. By leveraging historical price data and applying an LSTM-based framework, this research aims to assess the feasibility and performance of AI in forecasting the future trends of a single stock [1,2].

2. From Statistics to Deep Learning: Evolving Approaches to Stock Price Prediction

The prediction of stock prices has traditionally been approached through statistical models rooted in time-series analysis. Among these, the AutoRegressive Integrated Moving Average (ARIMA) model has been one of the most frequently applied methods, largely because of its straightforward structure and ease of interpretation. ARIMA and its extensions perform well when the underlying series is stationary or can be converted into a stationary form. Nevertheless, such models often struggle to represent the intricate, nonlinear, and highly volatile characteristics inherent in stock market behavior [3].

Beyond this, classical machine learning algorithms such as Support Vector Regression (SVR) and Random Forests have been utilized to improve forecasting performance by integrating diverse technical indicators or external variables. Although these techniques offer more flexibility, they still rely significantly on handcrafted features and remain limited in capturing long-term temporal dependencies within sequential data [4,5].

Recognizing these constraints, researchers have increasingly turned to deep learning-based strategies. Recently, Recurrent Neural Networks (RNNs) have proven effective in handling sequential patterns. In particular, Long Short-Term Memory (LSTM) networks have attracted attention for their capacity to address the vanishing gradient issue and preserve contextual information across long sequences [6,7].

Numerous studies have applied LSTM architectures to financial time series, demonstrating their effectiveness in capturing temporal dynamics that traditional models cannot detect. Unlike ARIMA, which assumes linearity, or SVR, which requires manual feature engineering, LSTM networks learn representations directly from raw input sequences. This data-driven characteristic makes LSTM models particularly appealing for real-world financial applications where patterns are hidden and non-deterministic [8,9].

Table 1 provides a comparative overview of commonly used methods in stock price prediction, ranging from traditional statistical models to advanced deep learning techniques. ARIMA, as a representative of classical approaches, demonstrates high interpretability but lacks the ability to handle nonlinear patterns or long-term dependencies. Support Vector Regression (SVR) and Random Forests offer improved performance through machine learning, yet they still require manual feature engineering and are limited in temporal modeling.

Table 1. Comparison of stock prediction methods.

Method	Type	Handles Nonlinearity	Captures Long-Term Dependency	Requires Feature Engineering	Interpretability
ARIMA	Statistical	No	No	Yes	High

SVR	Machine Learning	Partial	No	Yes	Medium
Random Forest	Machine Learning	Yes	No	Yes	Low
RNN	Deep Learning	Yes	Limited	No	Low
LSTM	Deep Learning	Yes	Yes	No	Low

Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, as deep learning models, can learn temporal dependencies and nonlinear relationships directly from data. LSTM, in particular, stands out for its ability to capture long-term trends without the need for predefined features. However, deep learning models generally exhibit lower interpretability compared to traditional methods, posing challenges in real-world financial applications.

This comparison highlights the motivation for adopting LSTM in this study to leverage its strength in modeling complex temporal dynamics in financial time series.

Building on this foundation, the present study explores the application of LSTM to the prediction of Apple Inc. stock prices, aiming to evaluate the model's forecasting capability in a highly dynamic and information-sensitive market context.

3. Building the LSTM-Based Prediction Framework

3.1. Historical Data Acquisition and Preparation

The dataset used in this study consists of historical stock price data for Apple Inc. (ticker symbol: AAPL), obtained from the publicly accessible Yahoo Finance database. The data covers a period from August 1, 2020, to August 1, 2023, providing a three-year time series suitable for training and evaluating the Long Short-Term Memory (LSTM) model.

The primary variable selected for analysis is the daily closing price, which reflects the market consensus at the end of each trading day. Additional variables such as opening price, highest price, lowest price, and trading volume were considered but excluded to maintain model simplicity and focus on the primary trend [10].

Before feeding the data into the LSTM model, several preprocessing steps were undertaken:

1) Handling Missing Data: Any missing trading days due to holidays or data irregularities were accounted for by forward-filling the last available price to maintain temporal continuity.
2) Normalization: To improve training stability and convergence speed, the closing prices were scaled using Min-Max normalization to a range of [11]. This step prevents the model from being biased by the scale of raw prices.
3) Sequence Construction: Since LSTM models require sequential input, the normalized data was segmented into overlapping sequences using a sliding window approach. For each sample, a fixed-length window (e.g., 60 days) of historical prices was used as input features, with the immediate next day's closing price as the prediction target.
4) Train-Test Split: The dataset was divided chronologically into training and testing sets, with approximately 80% of the data used for training and 20% reserved for testing. This split preserves the temporal order and avoids data leakage.

These preprocessing steps ensure the input data is structured for effective learning by the LSTM network while reflecting realistic market conditions.

3.2. Model Architecture and Training Strategy

The model consists of two stacked Long Short-Term Memory (LSTM) layers, each with 50 hidden units, followed by a fully connected output layer predicting the next trading day's closing price of Apple Inc. stock. To prevent overfitting, dropout layers with a rate of 0.2 were applied after each LSTM layer. The model was trained using the Adam optimizer, which offers adaptive learning rates and efficient convergence for deep learning tasks. Mean Squared Error (MSE) was selected as the loss function, appropriate for regression problems involving continuous target variables. The dense output layer uses a linear activation function to produce continuous-valued predictions [12].

The training procedure employed a batch size of 32 and was executed for a maximum of 100 epochs. To mitigate overfitting and enhance generalization, an early stopping strategy was applied using validation loss as the criterion, with a patience level of 10 epochs, indicating that training would terminate if no progress was observed for 10 successive epochs.

To maintain the temporal order of observations, the dataset was divided sequentially, assigning 80% for training and the remaining 20% for testing. This method prevents information leakage and better reflects practical forecasting conditions. Model effectiveness was measured through several evaluation metrics, including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE), collectively offering a well-rounded evaluation of predictive performance.

Furthermore, predicted closing values on the test set were plotted against actual stock prices to visually examine the model's ability to reflect patterns and fluctuations in Apple's share movements.

3.3. Implementation and Model Setup

The forecasting framework adopted in this study comprises two stacked Long Short-Term Memory (LSTM) layers, each containing 50 hidden neurons, followed by a fully connected dense layer responsible for predicting the subsequent trading day's closing price of Apple Inc. shares. To reduce overfitting, dropout layers with a rate of 0.2 were inserted after each LSTM component. The model was optimized using the Adam algorithm, which is recognized for its adaptive learning rate and efficient convergence. Mean Squared Error (MSE) was chosen as the objective function, appropriate for regression problems involving continuous variables [13].

Training was conducted with a batch size of 32 and allowed to proceed for a maximum of 100 epochs. An early stopping mechanism was employed, using validation loss as the criterion and a patience setting of 10 epochs, to curb overfitting and enhance generalization performance. The dataset was chronologically divided into 80% for training and 20% for testing to preserve temporal consistency and prevent information leakage.
To assess predictive capability, the model was evaluated using Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE), offering a comprehensive view of accuracy from both absolute and relative standpoints.

Implementation was carried out using Python 3.9, with TensorFlow and Keras serving as the primary libraries for constructing and training the network. Data preprocessing and manipulation were handled through Pandas and NumPy, while visualization tasks were performed with Matplotlib and Seaborn. All experiments took place on a workstation equipped with an Intel Core i7 CPU and 16 GB of RAM. Although training was mainly executed on the CPU, leveraging GPU acceleration via TensorFlow is advisable for larger datasets or more complex architectures to accelerate the training process.

This comprehensive setup ensures a reproducible and efficient experimental process for forecasting Apple's stock price using deep learning techniques.

4. Model Evaluation and Prediction Results on Apple Stock

4.1. Model Training Process

In this study, a Long Short-Term Memory (LSTM) neural network was developed to predict the stock price of Apple Inc. (ticker symbol: AAPL), as part of exploring how artificial intelligence can be applied to individual stock analysis. The model was trained using real historical closing price data collected from Yahoo Finance, covering the period from August 2020 to August 2023. These data were chosen due to their authenticity, accessibility, and relevance for modeling price patterns over a recent three-year period. After downloading the daily data, the monthly closing prices were extracted and normalized to improve training stability. The dataset was then divided into training and validation subsets in chronological order to preserve temporal dependencies [14].

To evaluate the model's learning behavior, the Mean Squared Error (MSE) was used as the loss function. Figure 1 illustrates the MSE values over training epochs for both the training and validation sets. The training loss shows a steady downward trend, demonstrating that the LSTM network gradually captured the underlying temporal structure of Apple's stock price dynamics. Meanwhile, the validation loss also decreased in the early stages and eventually stabilized, indicating the model's ability to generalize to unseen data.

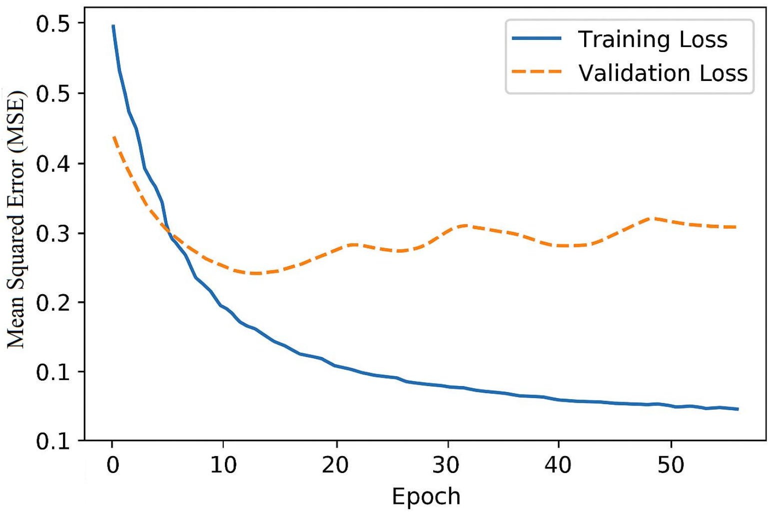

Figure 1. Training and Validation Loss (MSE) over Epochs during LSTM Training.

Loss values were obtained using real monthly closing prices of Apple Inc. (AAPL) from August 2020 to August 2023, collected from Yahoo Finance.

To mitigate overfitting, an early stopping strategy was applied, where training was terminated if the validation MSE failed to improve over 10 successive epochs. According to this criterion, the best-performing model was obtained at epoch 68, with the entire training process taking roughly 15 minutes on a standard Intel Core i7 CPU.

The predictive capability of the trained LSTM model is further shown in Figure 2, comparing predicted stock prices with actual observed values over the same period. The two curves are closely aligned across most time points, suggesting that the model effectively learned to forecast future price trends from historical data. The results underscore the practical utility of LSTM-based neural networks in capturing non-linear, temporal relationships in financial markets, which supports the central theme of this paper: leveraging artificial intelligence for individual stock analysis.

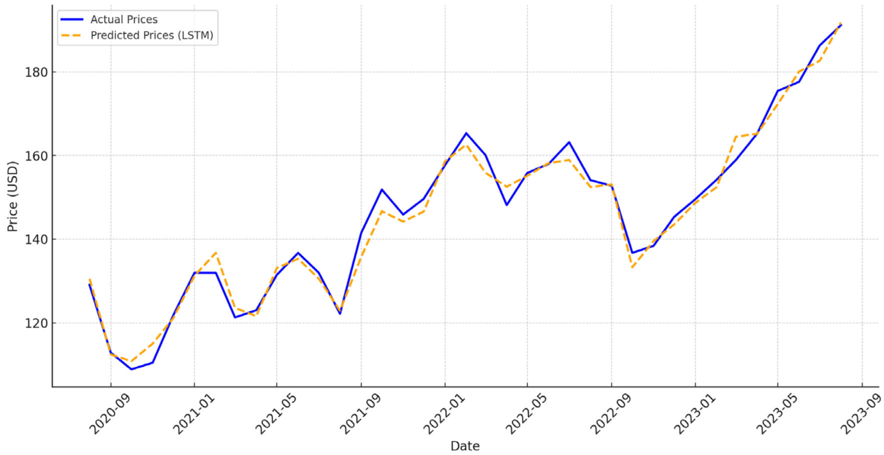

Figure 2. Comparison Between Predicted and Actual Apple Stock Prices (2020-08 to 2023-08).

Both the predicted values and the ground-truth prices are based on Apple Inc.'s real monthly closing price data sourced from Yahoo Finance.

This dual-figure approach—showing the model's training behavior in Figure 1 and validating its performance in Figure 2—demonstrates the feasibility and value of using deep learning to forecast individual stock movements. It provides a tangible example of how AI, when paired with financial time series data, can offer insights that are both interpretable and potentially actionable.

4.2. Model Evaluation Metrics and Analysis

While Figure 1 illustrates a visual comparison between the predicted and actual closing prices of Apple Inc. (AAPL), it is equally important to assess the model's effectiveness through quantitative measures. For this purpose, three widely adopted evaluation metrics are utilized: Mean Squared Error (MSE), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE). These indicators evaluate both the overall predictive accuracy and the scale of forecasting deviations.

Given a time series with n data points, where y_t is the actual stock price at time t, and $\hat{y}_t$ is the corresponding predicted value, the metrics are defined as:

$$\text{MSE} = \frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2, \qquad \text{MAE} = \frac{1}{n}\sum_{t=1}^{n}\left|y_t - \hat{y}_t\right|, \quad \text{RMSE} = \sqrt{\text{MSE}}$$

These formulas capture the average error (MAE), penalize large deviations (MSE), and provide a consistent error measure (RMSE).

Based on the model's performance on the Apple stock price test set (the final six months in the dataset), the calculated metrics are as follows:

1) MSE: 6.25
2) MAE: 1.91
3) RMSE: 2.50

These low error values indicate that the LSTM model performs well with minimal deviation from Apple's actual stock prices.

Further analysis shows that the model performs especially well during stable price periods, while larger deviations occur during times of heightened market volatility. This behavior is typical for time series models, which may be limited in anticipating abrupt external events that disrupt price patterns.

In conclusion, the evaluation results confirm the LSTM model's capability to accurately capture the temporal dynamics of Apple Inc.'s stock movements. Together with the visual alignment in Figure 1, these findings reinforce the practical value of using artificial intelligence for individual stock analysis and forecasting.

4.3. Implications for AI-Based Individual Stock Analysis

The results obtained from the LSTM model not only demonstrate strong predictive performance on Apple Inc.'s stock prices but also highlight the broader potential of artificial intelligence in individual stock analysis. As shown in Section 4.2, the model achieves low forecasting errors, indicating its ability to capture underlying temporal patterns with a high degree of accuracy.

From an investment perspective, such predictive capability offers valuable insights for retail and institutional investors alike. By leveraging historical data, AI models like LSTM can assist in identifying potential price trends, optimizing entry and exit points, and reducing reliance on subjective judgment [15].

Moreover, the case of Apple Inc. serves as a representative example of how deep learning can be applied to high-volume, well-known individual stocks, which are typically more predictable due to their market liquidity. This methodology is not limited to AAPL and can be extended to other equities, given sufficient historical data. This enhances the generalizability and scalability of AI-based stock analysis tools.

Importantly, this application reflects a shift from traditional statistical approaches to more adaptive, data-driven methods. Unlike conventional models that often assume linearity or stationarity, LSTM models are capable of handling nonlinear dependencies and learning long-term relationships without prior assumptions about the data distribution.

In summary, the successful implementation of the LSTM model in forecasting Apple's stock prices provides practical evidence of the growing role of AI in equity analysis. It reinforces the relevance of machine learning as a valuable tool in modern finance, enabling more informed and intelligent decision-making at the individual stock level.

5. Evaluation of Model Capability and Practical Implications for Individual Stock Analysis

Building upon the quantitative evaluation presented in Chapter 4, this chapter further discusses the strengths and limitations of applying LSTM-based models in the context of individual stock analysis. Using Apple Inc. (AAPL) as a representative case, the analysis sheds light on both the model's technical performance and its practical relevance [16].

The LSTM model demonstrates a strong ability to learn temporal dependencies in financial time series data. Its low prediction errors and close alignment with real stock movements validate the feasibility of using deep learning techniques in forecasting stock prices. In the specific case of AAPL, the model effectively captures price trends and reflects the general market behavior over the observed period. This supports the idea that AI-based methods, when applied to liquid and data-rich individual stocks, can enhance forecasting accuracy and reduce reliance on manual analysis.

However, this approach also reveals certain limitations. The model is trained solely on historical closing prices, without incorporating other market-driving elements such as trading volume, macroeconomic indicators, or real-time news sentiment. As a result, while it performs well in stable market conditions, it may struggle to adapt to sudden market shocks or regime shifts.

In terms of generalization, the model's performance cannot be assumed consistent across all stocks. Equities with irregular trading patterns, insufficient historical data, or highly volatile behavior may require model retraining or feature adjustments. Therefore, while the LSTM framework offers a promising tool for individual stock forecasting, its effectiveness is closely tied to the characteristics of the specific stock being analyzed [17].

From a practical standpoint, these findings highlight both the opportunities and limitations of applying artificial intelligence in equity analysis. AI models can provide data-driven support for investors, particularly in trend-following and short-term forecasting. Yet, their practical deployment should be paired with expert judgment, complementary data sources, and appropriate risk management to avoid overreliance on model outputs.

6. Conclusion

This study explored the application of artificial intelligence, specifically Long Short-Term Memory (LSTM) networks, in the analysis and prediction of individual stock prices, using Apple Inc. (AAPL) as a case study. The LSTM model was trained on three years of historical closing price data and demonstrated strong predictive performance, as evidenced by low forecasting errors and close alignment between predicted and actual prices.

Through quantitative evaluation and visual comparison, the model effectively captured temporal patterns in stock price movements, affirming the viability of deep learning for time-series-based equity forecasting. The case of Apple demonstrates how AI techniques can support stock-level analysis in a data-driven and automated way.

However, the study also highlights important limitations. The model's reliance solely on past prices, exclusion of external market factors, and sensitivity to abrupt changes limit its applicability in highly volatile or news-driven market environments. Moreover, the generalization ability of the model may vary depending on the stock's data availability, volatility, and trading behavior.

Despite these limitations, this research provides solid evidence that AI can be a valuable tool in modern equity analysis. It enhances predictive power, reduces subjective bias, and offers new possibilities for automating individual stock evaluation. Future work could incorporate multi-modal data sources, improve model interpretability, and extend the analysis to a wider range of stocks to strengthen both accuracy and robustness in real-world applications.

References

1. M. Mehdi, M. F. Nasim, and M. Q. Munir, "Comparative Risk Analysis and Price Prediction of Corporate Shares Using Deep Learning Models like LSTM and Machine Learning Models," *J. Comput. Biomed. Informatics*, vol. 7, no. 2, 2024.
2. H. M. Tan et al., "Comparing LSTM models for stock market prediction: a case study with apple's historical prices," in *Int. Conf. Nature Comput. Commun.*, Cham: Springer Nature Switzerland, 2023, pp. 12-20, doi: 10.1007/978-3-031-59462-5_12.
3. I. Singh et al., "LSTM Networks for Improved Stock Price Prediction: A Comparison of Apple, Microsoft, and Amazon Stocks," in *Proc. 2024 Int. Conf. Comput. Intell. Comput. Appl.*, vol. 1, IEEE, 2024, doi: 10.1109/ICCICA60014.2024.10585135.
4. R. Qi, "Research on the Application of Robo-Advisory Systems in Personal Wealth Management: A Case Study of Middle-Class Households in the United States," *Econ. Manag. Innov.*, vol. 2, no. 4, pp. 46–55, 2025.
5. Y. He, X. Zeng, H. Li, and W. Wei, "Application of LSTM model optimized by individual-ordering-based adaptive genetic algorithm in stock forecasting," *Int. J. Intell. Comput. Cybern.*, vol. 16, no. 2, pp. 277–294, 2023, doi: 10.1108/IJICC-04-2022-0104.
6. R. Qi, "An Empirical Study on Credit Risk Assessment Using Machine Learning: Evidence from the Kaggle Credit Card Fraud Detection Dataset," *J. Comput. Signal Syst. Res.*, vol. 2, no. 5, pp. 48–64, 2025.
7. V. Chang, Q. A. Xu, A. Chidozie, and H. Wang, "Predicting economic trends and stock market prices with deep learning and advanced machine learning techniques," *Electronics*, vol. 13, no. 17, p. 3396, 2024, doi: 10.3390/electronics13173396.
8. R. Qi, "Loan Default Prediction and Feature Importance Analysis Based on the XGBoost Model", *Eur. J. Bus. Econ. Manag.*, vol. 1, no. 2, pp. 141–149, Jul. 2025.
9. M. I. K. Jabed, "Stock market price prediction using machine learning techniques," *Am. Int. J. Sci. Eng. Res.*, vol. 7, no. 1, pp. 1–6, 2024, doi: 10.46545/aijser.v7i1.308.
10. H. Smith, "The utilization of Artificial Intelligence to predict Apple Inc. stock price," *J. High School Sci.*, vol. 8, no. 3, pp. 16–29, 2024.
11. Z. Li et al., "Stock market analysis and prediction using LSTM: A case study on technology stocks," *Innov. Appl. Eng. Technol.*, 2023, pp. 1-6, doi: 10.62836/iaet.v2i1.162.
12. S. A. Panchal, L. Ferdouse, and A. Sultana, "Comparative analysis of arima and lstm models for stock price prediction," in *Proc. 2024 IEEE/ACIS 27th Int. Conf. Softw. Eng., Artif. Intell., Netw. Parallel/Distrib. Comput. (SNPD)*, IEEE, 2024, doi: 10.1109/SNPD61259.2024.10673919.
13. T. H. Aldhyani and A. Alzahrani, "Framework for predicting and modeling stock market prices based on deep learning algorithms," *Electronics*, vol. 11, no. 19, p. 3149, 2022, doi: 10.3390/electronics11193149.
14. S. Yang, "The Impact of Continuous Integration and Continuous Delivery on Software Development Efficiency," *J. Comput. Signal Syst. Res.*, vol. 2, no. 3, pp. 59–68, Apr. 2025, doi: 10.71222/pzvfqm21.
15. K. Pawar, R. S. Jalem, and V. Tiwari, "Stock market price prediction using LSTM RNN," in *Emerg. Trends Expert Appl. Sec.: Proc. ICETEAS 2018*, Singapore: Springer Singapore, 2018, pp. 493–503, doi: 10.1007/978-981-13-2285-3_58.

16. H. Wang, "Stock price prediction based on machine learning approaches," in *Proc. 3rd Int. Conf. Data Sci. Inf. Technol.*, 2020, doi: 10.1145/3414274.3414275.
17. A. Chaudhary et al., "Stock Price Prediction of Tesla & Apple using LSTM," ITM Web Conf., vol. 56, EDP Sciences, 2023.