 e-ISSN: 2686-2573
DOI: 10.21512/emacsjournal.v7i3.14326

Comparative Analysis of CNN, LSTM, and CNN–LSTM for Indonesian Stock Prediction

Setiawan Joddy

Computer Science Program, Computer Science Department,
School of Computer Science,
Bina Nusantara University,
Jakarta, Indonesia 11480
setiawan.joddy@binus.ac.id

Correspondence: setiawan.joddy@binus.ac.id

Abstract *– Predicting stock market movements remains a challenging task due to the nonlinear, volatile, and dynamic characteristics of financial time series. Traditional statistical methods often fall short in capturing these complexities, motivating the use of deep learning approaches that can learn hierarchical representations and temporal dependencies from sequential data. While deep learning has been widely adopted in developed markets, research focusing on emerging markets such as Indonesia is still relatively limited. Addressing this gap, this study conducts a comparative analysis of three deep learning models—Convolutional Neural Networks (CNN), Long Short-Term Memory networks (LSTM), and a hybrid CNN–LSTM—on five randomly selected constituents of the IDX30 index. The dataset we utilized spans from January 2020 to December 2024, providing a comprehensive perspective on recent stock price dynamics. Daily OHLCV (Open, High, Low, Close, Volume) data were collected and preprocessed using normalization and a sliding-window approach to transform the series into supervised learning format. All models were trained under consistent hyperparameter settings to ensure comparability. Results demonstrate that LSTM outperformed the other models, achieving the lowest RMSE (0.0222 ± 0.0030), lowest MAE (0.0172 ± 0.0015), and the highest* R^2 *(0.889 ± 0.068). The Hybrid CNN–LSTM ranked second, outperforming CNN but not surpassing LSTM, while CNN consistently yielded weaker results. These findings confirm that LSTM networks are particularly effective for stock prediction in the Indonesian context, while hybrid models offer complementary benefits by balancing local feature extraction with long-term temporal modeling.*

Keywords: *CNN; LSTM; Deep Learning; Indonesian Stock Exchange (IDX30); Stock Market Prediction*

I. INTRODUCTION

Stock market forecasting has been an essential problem in computational finance, as accurate predictions will have significant investment decisions and risk management strategies. Traditional statistical methods such as AutoRegressive Integrated Moving Average (ARIMA) and Generalized AutoRegressive Conditional Heteroskedasticity (GARCH), have been widely used to model financial time series (Dadhich et al., 2021; Marisetty, 2024; Srivastava et al., 2024). But these approaches struggle to capture the complex, non-linear, and dynamic behavior of stock price movements. Deep learning approaches have emerged over the past few years as potential alternatives as they can handle sequential data and tap into inderent patterns in large data (Shah et al., 2022).

Among numerous deep learning methods, Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks have received significant attention. CNNs, originally developed for computer vision applications, have been found very effective for the extraction of local temporal patterns from time series data by utilizing the convolutional filters, enabling useful learning of short-term dependencies (Jiang et al., 2023). In forecast finance, 1-D CNNs are extensively employed to handle sequences of OHLCV (Open, High, Low, Close) data in order for the model to acquire short-term patterns of the market (Gupta et al., 2023; Mezghani & Abbes, 2023).

On the other hand, LSTM also have been specifically developed to learn long-run temporal patterns and circumvent the vanishing gradient problem of ordinary RNN (Fadziso, 2020; Noh, 2021; Xiao & Sun, 2021). LSTM are usually applied in financial prediction tasks since they can preserve memory of lagged relationships with long time horizons and hence represent a suitable candidate for depicting the sequentiality of stock market behavior (Aswini et al., 2024; Choudhury et al., 2023; Laxmi Narayan & Balaji, 2025).

Recent findings shows that hybrid CNN–LSTM models to capitalize on the strengths of the individual models. Feature extractors, CNN layers learn local temporal information, while LSTM layers learn the long-range temporal dependencies. The unique combination of the two has been proven to improve predictive performance in financial forecasting over standalone CNN or LSTM models (Joshi et al., 2025; Lu et al., 2020; Zhao et al., 2024).

Despite the widespread use of deep learning for stock market prediction in the globe, research on newer markets, including Indonesia, is comparatively fewer. Indonesian Stock Exchange (IDX), or the IDX30 index, represents a significant portion of the country's capital market turnover, but there is not much research on deep learning-based forecasts on stocks in Indonesia. Also, global and local markets have experienced heightened volatility in recent years, which justifies the application of robust forecasting models with the ability to respond to changing financial environments (Hong et al., 2021).

These gaps are filled in this study by conducting comparative analysis between CNN, LSTM, and CNN-LSTM hybrid models for stock price forecasting using data for five randomly selected companies on the IDX30. The data is from January 2020 to December 2024 and covers a complete representation of stock price movements under different market conditions. There are three aims of this research: (1) to evaluate the predictive effectiveness of CNN, LSTM, and CNN–LSTM models in forecasting Indonesian stock prices, (2) to compare their performance using a consistent dataset, implementation framework, and evaluation metrics (RMSE, MAE, and R^2), and (3) to provide insights into the applicability of deep learning models for stock forecasting purposes in the context of emerging markets.

II. METHODS

2.1. Dataset

The dataset used for this study consists of daily stock price data from five randomly selected companies listed on the IDX30 index, which represents some of the most liquid and large-cap stocks on the Indonesian Stock Exchange. The observation period spans from January 2020 to December 2024, providing a comprehensive five-year view of trading activity under different market conditions. Data were obtained using the yfinance Python library, which provides direct access to historical market data from Yahoo Finance.

The use of yfinance ensures reproducibility and facilitates efficient extraction of OHLCV attributes. For each trading day, five standard features were collected: Open, High, Low, Close, and Volume (OHLCV). These features are widely used in stock forecasting tasks because they reflect both price dynamics and market participation, making them suitable inputs for predictive modeling (Bhardwaj & Singh, 2023; Mane et al., 2025). By focusing on IDX30 constituents, the dataset captures actively traded stocks that are representative of the Indonesian market structure.

2.2. Data Preprocessing

Before training, the raw stock data underwent a series of preprocessing steps to ensure the consistency. Missing values, which can arise due to non-trading days or data reporting issues were addressed using forward filling, a method that substitutes missing entries with the most recent available value.

Next, all features were normalized into the range [0,1] using MinMaxScaler, ensuring comparability across variables with different magnitudes. This step prevents features with larger scales, such as trading volume, from disproportionately influencing the learning process. To reframe the problem into a supervised learning task, a sliding window approach was applied, where the past N trading days were used as inputs to predict the next day’s closing price. We use the look-back windows of 45 days. Finally, the dataset was divided into training and testing sets using by 80:20 chronological split to reflect real-world forecasting conditions.

2.3. Model Architectures

This study evaluates three deep learning architectures: CNN, LSTM, and CNN-LSTM. All architectures were implemented in TensorFlow with the same training configuration to ensure comparability. Across all models, the following hyperparameters were applied: a learning rate of 0.001, batch size of 32, and a maximum of 100 epochs with early stopping (patience of 10). The overall research workflow is summarized in Figure 1, which illustrates the sequential process from dataset collection to data preprocessing, model training, and performance evaluation. After preprocessing, the data were used to train CNN, LSTM, and CNN–LSTM architectures, each designed to capture different aspects of stock market dynamics. The final stage of the workflow involved evaluating the models with RMSE, MAE, and R^2 metrics.

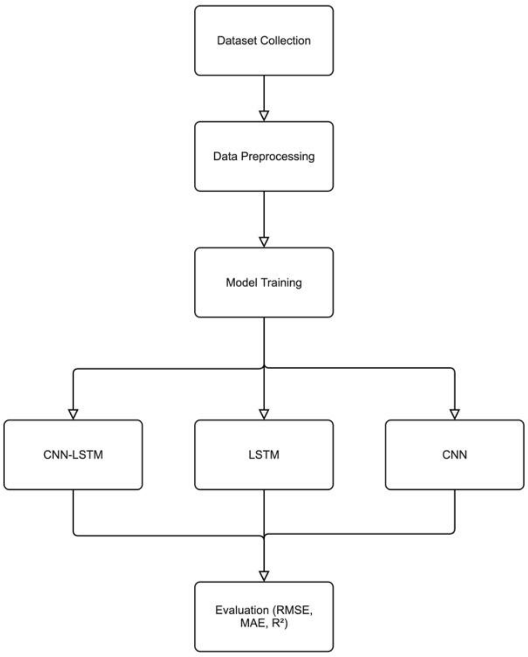

Figure 1. Model Architectures

The CNN model employs one-dimensional convolutional layers applied to OHLCV input sequences. Each convolutional layer uses 32 filters with a kernel size of 2 and ReLU activation to extract local temporal dependencies such as price momentum and volatility shifts. Pooling layers are added to reduce dimensionality and mitigate overfitting, and the resulting feature maps are fed into dense layers for regression output.

For the LSTM model, it consists of stacked recurrent layers with 64 LSTM units each, followed by fully connected dense layers. The gating mechanism of LSTM allows the model to retain relevant information while filtering noise, making it suitable for modeling longer-term stock market dynamics.

Last, the hybrid CNN-LSTM model integrates convolutional and recurrent layers to leverage both short-term and long-term temporal features. In this design, one-dimensional convolutional layers with 32 filters act as feature extractors to capture local temporal patterns and reduce noise. The extracted features are then passed to LSTM layers with 64 units, which learn dependencies over longer horizons. Dense layers are subsequently applied to produce the final regression output.

2.4. Implementation

The models trained in this study were executed in Python 3.12 using the TensorFlow library. Training was performed in a GPU-based setup to reduce computation time. Hyperparameters such as learning rate, batch size, number of convolution filters, number of LSTM units, and the look-back window size were tuned experimentally through grid search. To prevent overfitting and improve the stability of training, the early stopping method was utilized based on validation loss, which ended training after no improvement was observed after a specified number of epochs.

2.5. Evaluation Metrics

All models were evaluated for their predictive validity using three popular measures of regression: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and the measure of the coefficient of determination (R^2). These are utilized as they provide complementary perspectives towards the prediction accuracy and model fit.

Root Mean Square Error (RMSE) measures the magnitude of prediction errors by computing the square root of the average of squared absolute differences between predicted and actual values. RMSE is more weighted on bigger errors and, as such, is sensitive to outliers.

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

Mean Absolute Error (MAE) computes the mean of absolute absolute differences between actual and predicted values. MAE is equally weighted on all errors as opposed to RMSE and, as such, is less sensitive to outliers.

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

Coefficient of Determination (R^2) measures the extent to which the estimated values capture the actual values. R^2 ranges from negative infinity to 1, and the higher values indicate more explanatory power.

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

Together, RMSE, MAE, and R^2 provide a comprehensive evaluation of model performance: RMSE emphasizes large errors, MAE captures average error magnitude, and R^2 measures the proportion of variance explained by the model.

III. RESULTS AND DISCUSSION

3.1. Experimental Results

The prediction results of the three models were evaluated using daily OHLCV data from five IDX30 companies: ANTM.JK, BBCA.JK, CPIN.JK, TLKM.JK, and UNVR.JK. Each model was trained using the same hyperparameters (learning rate = 0.001, batch size = 32, convolution filters = 32, LSTM units = 64, and epochs = 100). This uniform setup ensures that the performance differences are attributable to the model architectures rather than variations in training configuration.

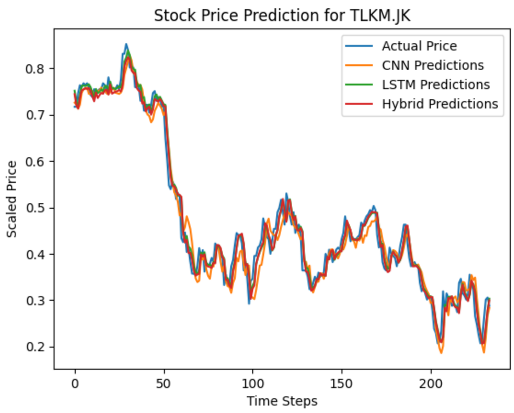

Figure 2. Comparisons between Actual, CNN, LSTM, and CNN-LSTM on TLKM.JK

The results for each stock are reported in Table 1, which lists RMSE, MAE, and R^2 scores for all three models. These per-stock metrics provide insight into how model performance varies across different equities. For instance, TLKM.JK achieved consistently high R^2 values (>0.95) across all models, suggesting that its stock price patterns are relatively stable and predictable. Conversely, UNVR.JK exhibited notable disparities. The LSTM model achieved an excellent R^2 of 0.933 but both CNN and Hybrid models struggled with the score R^2 of 0.661 and 0.749, respectively. This findings highlight the sensitivity of models to stock-specific dynamics and volatility.

Table 1. Model Performance Across Five Stocks

Ticker	Model	RMSE	MAE	R^2
ANTM.JK	CNN	0.0214	0.0171	0.757
	LSTM	0.0195	0.0155	0.798
	CNN-LSTM	**0.0175**	**0.0138**	**0.837**
BBCA.JK	CNN	0.0245	0.0194	0.842
	LSTM	**0.0218**	**0.0177**	**0.874**
	CNN-LSTM	0.0235	0.0183	0.854
CPIN.JK	CNN	0.0290	0.216	0.793
	LSTM	**0.0235**	**0.0171**	**0.865**
	CNN-LSTM	0.0249	0.0190	0.847
TLKM.JK	CNN	0.0343	0.0267	0.958
	LSTM	**0.0266**	**0.0195**	0.974
	CNN-LSTM	0.0263	**0.0195**	**0.975**
UNVR.JK	CNN	0.0441	0.0344	0.661
	LSTM	**0.0196**	**0.0163**	**0.933**
	CNN-LSTM	0.0379	0.0284	0.749

While stock-level results illustrate differences in predictive difficulty, they do not fully convey overall performance trends. To address this, the results were aggregated across the five stocks, and the mean ± standard deviation of each metric was computed. Table 2 summarizes these aggregated outcomes.

Table 2. Average Predictive Performance (mean ± std) Across Five Stocks

Model	RMSE (mean ± std)	MAE (mean ± std)	R^2 (mean ± std)
CNN	0.0306 ± 0.0090	0.0238 ± 0.0069	0.802 ± 0.109
LSTM	**0.0222 ± 0.0030**	**0.0172 ± 0.0015**	**0.889 ± 0.068**
CNN-LSTM	0.0260 ± 0.0075	0.0198 ± 0.0053	0.853 ± 0.081

3.2. Comparative Analysis of Models

The aggregated metrics in Table 2 confirmed that LSTM consistently outperformed CNN and Hybrid CNN-LSTM across all three evaluation metrics. With the lowest RMSE (0.0222 ± 0.0030) and MAE (0.0172 ± 0.0015), and the highest R^2 (0.889 ± 0.068), LSTM demonstrated not only superior accuracy but also more stable performance across different stocks, as evidenced by its relatively small standard deviations.

The Hybrid CNN-LSTM model achieved intermediate results. It performed better than CNN on average, showing improvements in both RMSE and R^2. However, it did not surpass LSTM. This finding suggests that while convolutional layers can aid in feature extraction and noise reduction, their added complexity does not always lead to performance gains when LSTM alone is already capable of capturing long-term temporal dependencies.

The CNN model recorded the weakest results, with the highest errors (RMSE = 0.0306 ± 0.0090; MAE = 0.0238 ± 0.0069) and lowest explanatory power (R^2 = 0.802 ± 0.109). Moreover, CNN exhibited higher variability across stocks, implying that it is less robust in adapting to different equity characteristics. These limitations highlight CNN's inability to capture the long-range dependencies essential for stock market forecasting, where historical sequences influence future price dynamics beyond local patterns.

3.3. Stock-Level Observations

Examining stock-level variations uncovers additional insights. For example, TLKM.JK differentiated itself with strong predictability, with R^2 values over 0.95 for every model as figured in the Figure 2. This suggests that certain Indonesian equities exhibit relatively smooth and stable price movements, making them easier to model. TLKM is among the most liquid and heavily traded stocks in the IDX30, which may explain why its patterns are captured effectively even by low-complexity models such as CNN. High liquidity tends to reduce irregular price swings, resulting in time series that are less noisy and more predictable.

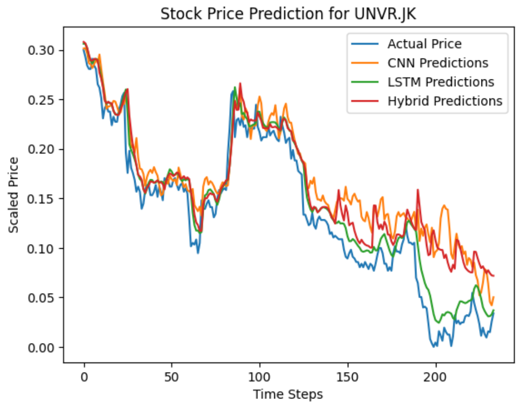

Figure 3. Comparisons between Actual, CNN, LSTM, and CNN-LSTM on UNVR.JK

In contrast, UNVR.JK posed a greater challenge, as shown in Figure 3. CNN and Hybrid CNN-LSTM models performed poorly, achieving low R^2 values (0.661 and 0.749, respectively), whereas LSTM achieved much better accuracy (R^2 = 0.933). This difference likely reflects the stock's higher volatility and less regular trading patterns. Consumer goods companies such as UNVR are often more exposed to shifts in domestic demand and external market pressures, which can lead to noisier signals that are harder for models to approximate. In this case, only the LSTM's ability to retain longer-term dependencies allowed it to capture underlying trends effectively.

Intermediate patterns were observed for other equities. ANTM.JK, a commodity-linked stock, showed relatively good performance under the Hybrid CNN-LSTM model, slightly outperforming LSTM. This may be due to the stock's sensitivity to external commodity price shocks, where the hybrid architecture's ability to combine short-term feature extraction with long-term dependency modeling becomes advantageous. Similarly, CPIN.JK exhibited moderate predictability as LSTM provided the strongest performance, but both CNN and Hybrid CNN-LSTM tracked general price movements with reasonable accuracy. These results illustrate that stocks with sectoral exposure to external shocks or irregular demand cycles may benefit from hybrid designs, though long-term memory provided by LSTM remains critical.

Finally, BBCA.JK, one of Indonesia's largest banking institutions, displayed predictability levels comparable to TLKM.JK, with consistently high R^2 scores across all models. As a highly liquid and systemically important stock, BBCA's relatively stable trading patterns made it more predictable than equities exposed to commodity price swings or consumer demand volatility. This reinforces the observation that liquidity and trading stability strongly influence model performance.

Overall, these findings emphasize that model performance is influenced not only by architecture but also by stock-specific characteristics such as liquidity, volatility, and sectoral exposure. Therefore, model selection for financial forecasting should remain sensitive to the nature of each equity, as performance can vary significantly even among constituents of the same index.

3.4. Discussion

Overall, the findings demonstrate that the most appropriate model to apply in forecasting Indonesian stock prices is LSTM as it outperforms CNN and Hybrid CNN-LSTM models when it comes to accuracy and reliability. The mid-performance of Hybrid shows that the combination of CNN and LSTM can be advantageous in certain scenarios, particularly for stocks whose local short-term behaviors are significant. However, in this study, the additional convolutional layers failed to yield consistent results. This outcome points out that hybrid models should not necessarily be assumed to perform better than their architectures; their benefit may depend on dataset size, feature composition, and market conditions.

CNN models while providing decent results, lagged behind the rest of the models considerably. Its reliance on local pattern recognition limits it to handle the more complex dependencies present in financial information. Nevertheless, its efficiency compared to its performance makes it an appropriate selection for usage where speed and explainability are considered more valuable than maximum accuracy.

Adding both aggregated and stock-level analyses strengthens the argument for LSTM as the most reliable architecture in the case. Stock-level results indicate where models work or not work, but aggregated output provides statistical evidence that LSTM consistently works with relatively low variance. Finally, these findings demonstrate the potential and limitations of deep learning in emerging economies. While advanced models such as LSTMs greatly improve prediction performance, there remains a challenge with extremely volatile stocks.

IV. CONCLUSION

This study examined the predictive performance of three deep learning models—CNN, LSTM, and Hybrid CNN–LSTM—applied to five randomly selected constituents of the IDX30 index using daily OHLCV data started from January 2020 to December 2024. The models were trained with a standardized configuration to ensure comparability. Then, the performance of each model was assessed using RMSE, MAE, and R^2.

The results of this study shows that LSTM consistently outperformed both CNN and Hybrid CNN-LSTM, achieving the lowest prediction errors on both RMSE = 0.0222 ± 0.0030 and MAE = 0.0172 ± 0.0015 and the highest R^2 score of 0.889 ± 0.068. The Hybrid model ranked second, performed better than CNN but not surpassing the LSTM, indicating that while CNN layers can enhance feature extraction, the primary driver of accuracy in this context is the LSTM's ability to capture long-term temporal dependencies. CNN, although computationally efficient, exhibited the weakest performance overall and higher variability across stocks, suggesting its limitations for financial time series forecasting in emerging markets.

At the stock level, performance varied across equities. For example, TLKM.JK, displayed strong predictability with high R^2 scores across all models, while others, such as UNVR.JK, posed greater challenges, where only LSTM achieved reliable accuracy. These findings highlight that model performance can depend not only on the architecture but also on stock-specific dynamics.

In conclusion, LSTM provides the most robust and reliable framework for predicting stock prices in the Indonesian market, with Hybrid CNN-LSTM offering competitive but not superior performance. This study contributes to the limited literature on deep learning in emerging financial markets and provides evidence-based insights for practitioners and researchers. Future research could extend this study by broadening the dataset to include more IDX30 constituents or other Indonesian stock indices, which would help assess whether the findings generalize across a wider range of equities. Another promising direction is the incorporation of additional features, such as technical indicators or macroeconomic variables to provide richer contextual information for forecasting. In terms of methodology, exploring advanced architectures such as attention-based models or Transformers may capture complex temporal patterns and long-range dependencies beyond what CNNs and LSTMs can offer.

ACKNOWLEDGMENT

This work did not receive any specific grant from funding agencies in the public, commercial, or non-profit sectors. The author acknowledges the use of generative AI tools, specifically OpenAI's ChatGPT, for assistance in language refinement and structural editing. All content was independently reviewed and validated to ensure accuracy and academic integrity.

REFERENCES

Aswini, J., S, D., Lakshmipriya, C., M, L. K., & R, S. S. (2024). Stock Market Forecasting Using LSTM. *2024 2nd World Conference on Communication & Computing (WCONF)*, 1–4. https://doi.org/10.1109/WCONF61366.2024.10692057

Bhardwaj, A., & Singh, U. P. (2023). An Empirical Study on: Time Series Forecasting of Amazon Stock Prices using Neural Networks LSTM and GAN models. *Proceedings of the 5th International Conference on Information Management & Machine Intelligence*, 1–10. https://doi.org/10.1145/3647444.3652440

Choudhury, S., Basak, S., Roy, S., & Das, A. K. (2023). Long-Short Term Memory Based Stock Market Analysis. *2023 7th International Conference on Electronics, Materials Engineering & Nano-Technology (IEMENTech)*, 1–5. https://doi.org/10.1109/IEMENTech60402.2023.10423465

Dadhich, M., Pahwa, M. S., Jain, V., & Doshi, R. (2021). Predictive Models for Stock Market Index Using Stochastic Time Series ARIMA Modeling in Emerging Economy. In G. Manik, S. Kalia, S. K. Sahoo, T. K. Sharma, & O. P. Verma (Eds.), *Advances in Mechanical Engineering* (pp. 281–290). Springer Singapore. https://doi.org/10.1007/978-981-16-0942-8_26

Fadziso, T. (2020). Overcoming the Vanishing Gradient Problem during Learning Recurrent Neural Nets (RNN). *Asian Journal of Applied Science and Engineering*, *9*(1), 197–208. https://doi.org/10.18034/ajase.v9i1.41

Gupta, A., Akansha, Joshi, K., Patel, M., & Pratap, Ms. V. (2023). Stock Prices Prediction Using Machine Learning. *2023 2nd International Conference for Innovation in Technology (INOCON)*, 1–7. https://doi.org/10.1109/INOCON57975.2023.10101226

Hong, H., Bian, Z., & Lee, C.-C. (2021). COVID-19 and instability of stock market performance: Evidence from the U.S. *Financial Innovation*, *7*(1), 12. https://doi.org/10.1186/s40854-021-00229-1

Jiang, M., Wang, K., Sun, Y., Chen, W., Xia, B., & Li, R. (2023). MLGN: Multi-scale local-global feature learning network for long-term series forecasting. *Machine Learning: Science and Technology*, *4*(4), 045059. https://doi.org/10.1088/2632-2153/ad1436

Joshi, S., Mahanthi, B. L., G, P., Pokkuluri, K. S., Ninawe, S. S., & Sahu, R. (2025). Integrating LSTM and CNN for Stock Market Prediction: A Dynamic Machine Learning Approach. *Journal of Artificial Intelligence and Technology*. https://doi.org/10.37965/jait.2025.0652

Laxmi Narayan, M. D., & Balaji, N. (2025). Stock Opening Price Prediction Using RNN and LSTM. In H. Sharma, A. Chakravorty, S. Hussain, & R. Kumari (Eds.), *Artificial Intelligence: Theory and Applications* (Vol. 5589, pp. 363–376). Springer Nature Singapore. https://doi.org/10.1007/978-981-96-1687-9_25

Lu, W., Li, J., Li, Y., Sun, A., & Wang, J. (2020). A CNN-LSTM-Based Model to Forecast Stock Prices. *Complexity*, *2020*, 1–10. https://doi.org/10.1155/2020/6622927

Mane, V., Patil, R., Pawar, S., Pujari, N., Jaggi, A., & Kakkar, P. (2025). Time Series Forecasting of Stock Prices Using Neural Networks LSTM and GAN. *2025 1st International Conference on AIML-Applications for Engineering & Technology (ICAET)*, 1–5. https://doi.org/10.1109/ICAET63349.2025.10932260

Marisetty, N. (2024). *Prediction of Popular Global Stock Indexes Volatility by Using ARCH/GARCH Models*. SSRN. https://doi.org/10.2139/ssrn.4904475

Mezghani, T., & Abbes, M. B. (2023). Forecast the Role of GCC Financial Stress on Oil Market and GCC Financial Markets Using Convolutional Neural Networks. *Asia-Pacific Financial Markets*, *30*(3), 505–530. https://doi.org/10.1007/s10690-022-09387-3

Noh, S.-H. (2021). Analysis of Gradient Vanishing of RNNs and Performance Comparison. *Information*, *12*(11), 442. https://doi.org/10.3390/info12110442

Shah, J., Vaidya, D., & Shah, M. (2022). A comprehensive review on multiple hybrid deep learning approaches for stock prediction. *Intelligent Systems with Applications*, *16*, 200111. https://doi.org/10.1016/j.iswa.2022.200111

Srivastava, A., Srivastava, A., Singh, Y. B., & Misra, M. K. (2024). Deep Learning Models for Stock Market Forecasting: GARCH, ARIMA, CNN, LSTM, RNN. In A. Chaturvedi, S. U. Hasan, B. K. Roy, & B. Tsaban (Eds.), *Cryptology and Network Security with Machine Learning* (Vol. 918, pp. 827–839). Springer Nature Singapore. https://doi.org/10.1007/978-981-97-0641-9_56

Xiao, C., & Sun, J. (2021). Recurrent Neural Networks (RNN). In C. Xiao & J. Sun, *Introduction to Deep Learning for Healthcare* (pp. 111–135). Springer International Publishing. https://doi.org/10.1007/978-3-030-82184-5_7

Zhao, Q., Hao, Y., & Li, X. (2024). Stock price prediction based on hybrid CNN-LSTM model. *Applied and Computational Engineering*, *104*(1), 110–115. https://doi.org/10.54254/2755-2721/104/20241065