Forecasting Stock Prices of Taiwan Semiconductor Manufacturing Company (TSMC) Using Recurrent Neural Networks: Evaluating Predictive Performance in a Volatile Market

Lai li-We[1*], Koundé Moanu[2]
[1]National Cheng Kung University, Taiwan
[2]Université Paris-Saclay, France
Email: lailiwe@gmail.com

Article Info

Article history:

Received: April 20, 2024
Revised: May 25, 2024
Accepted: July 10, 2024
Available Online: January 30, 2024

Keywords:

Stock Price Forecasting
Recurrent Neural Networks (RNN)
TSMC
Time Series Prediction
Financial Market Volatility

ABSTRACT

Accurate stock price prediction plays a critical role in guiding investment strategies, particularly in dynamic industries such as semiconductors, where price volatility is high. This study investigates the effectiveness of Recurrent Neural Networks (RNN) in predicting the stock prices of Taiwan Semiconductor Manufacturing Company (TSMC), a global leader in the semiconductor sector. Using daily closing price data from January 2020 to January 2023, the RNN model was developed and trained to forecast future stock prices. The data was preprocessed with feature scaling to ensure the stability of model training, and a sliding window approach was applied to capture temporal dependencies. The model's predictive performance was evaluated using Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE) as key metrics. The RNN achieved an RMSE of 9.87 and a MAPE of 5.90%, indicating that the model provides reasonable accuracy in forecasting stock prices with a moderate level of deviation from actual values. Visual analysis further demonstrated the model's capacity to capture general trends in the stock price movements, although challenges were noted in predicting highly volatile periods. The study highlights the potential of RNN in financial forecasting while suggesting future improvements, such as incorporating advanced models like Long Short-Term Memory (LSTM) or external factors to enhance predictions during market volatility. These findings offer valuable insights for investors and analysts seeking to leverage machine learning in stock price prediction, particularly in industries characterized by rapid technological advancements and price fluctuations.

1. INTRODUCTION

In the modern financial landscape, stock price forecasting plays a pivotal role in shaping investment decisions and risk management strategies [1]. Investors, hedge funds, and financial analysts rely on accurate predictions to optimize their portfolios and minimize losses [2]. However, predicting stock price movements remains a complex task due to the inherent volatility and randomness of financial markets[3], [4]. The development of machine learning techniques has opened new avenues for improving the accuracy of stock price forecasts, transforming how financial data is analyzed and utilized[5].

Machine learning models, particularly those designed for time series analysis, have demonstrated a significant potential in addressing the challenges of stock price forecasting [6]. Recurrent Neural Networks (RNN), in particular, stand out as a powerful tool for modeling sequential data [7]. Traditional machine learning methods, such as linear regression or even feedforward neural networks, struggle to capture the temporal dependencies that are crucial in financial time series [8]. RNNs, however, are designed to address this limitation by incorporating memory into the modeling process, making them highly suitable for dynamic, time-dependent data such as stock prices [9], [10].

The key innovation of RNNs lies in their architecture, which allows the model to "remember" information from previous time steps [11], [12]. This is achieved through a mechanism known as hidden states, where the output from one step is fed back into the network to influence the output at the next step [13]. By doing so, RNNs are capable of capturing long-term dependencies and patterns within time series data. This memory-driven approach is particularly beneficial in financial markets, where stock prices are influenced not only by recent events but also by trends that develop over time[14].

Despite the advantages of RNNs, stock price forecasting remains a challenging problem due to the complexity of market dynamics [15]. Stock prices are subject to a myriad of factors, including macroeconomic indicators, geopolitical developments, investor sentiment, and company-specific news. Accurately modeling these diverse influences requires sophisticated algorithms capable of identifying subtle patterns and correlations within the data [16]. RNNs, with their ability to process sequential and time-dependent information, offer a promising solution to this problem, yet their effectiveness in real-world scenarios requires further empirical validation [17].

Taiwan Semiconductor Manufacturing Company (TSMC) provides an ideal case study for evaluating the performance of RNNs in stock price forecasting. As the largest contract semiconductor manufacturer globally, TSMC plays a crucial role in the global technology supply chain. The company's stock is subject to significant volatility, driven by both internal factors such as production capacity and external factors like global demand for semiconductors [18]. Understanding and predicting TSMC’s stock price movements is of great interest to both institutional investors and industry analysts, given the company’s central role in the semiconductor market [19].

Given the complex nature of stock price movements and the high volatility associated with TSMC, traditional models may struggle to capture the intricate patterns present in the data [20]. In this study, we aim to leverage the advantages of RNNs to predict TSMC’s stock price, utilizing daily closing price data over a two-year period. The objective is to assess whether RNNs can effectively model the temporal dependencies in stock price data and provide more accurate predictions compared to traditional approaches.

Through this research, we contribute to the growing body of literature exploring machine learning applications in financial forecasting. By focusing on TSMC’s stock price and employing RNNs as the predictive model, this study aims to provide insights into the practical applications of deep learning models in stock price prediction. The findings will also offer valuable implications for investors and market analysts seeking to improve their decision-making processes using advanced predictive tools.

2. METHOD

2.1. Data Preparation

The stock price data for Taiwan Semiconductor Manufacturing Company (TSMC) was sourced from Yahoo Finance, covering the period from January 1, 2020, to January 1, 2023. The dataset used in this study includes the daily closing price (Close), which serves as the target variable for prediction. The following steps were undertaken to preprocess the data:

1. Feature Scaling: To improve model performance and accelerate convergence during training, the data was normalized using the MinMaxScaler. This scaler transforms the stock prices into the range [0, 1], which not only enhances the stability of model training but also mitigates the impact of large value discrepancies in the dataset.
2. Dataset Creation: A sliding window approach was employed to prepare the dataset, with a window size (time_step) of 60 days. This means that for each sample, the model utilizes the closing prices of the past 60 days as features (X) and the closing price of the following day as the target (y). The create_dataset function was used to generate these input-output pairs, allowing the model to learn patterns from historical stock prices and make predictions about future prices.

2.2. RNN Model Building

The Recurrent Neural Network (RNN) model was developed using the TensorFlow and Keras frameworks, both widely recognized platforms for implementing deep learning models. The architecture of the RNN is designed to capture temporal dependencies in the stock price data. The model consists of the following layers:

1. RNN Layers: Two RNN layers were implemented, each containing 50 units. The first RNN layer is configured with return_sequences=True, ensuring that the complete sequence from the previous time steps is passed to the subsequent layer. This provides the second RNN layer with more comprehensive temporal information. The second RNN layer is responsible for compressing the sequence data into a more compact representation, which is then passed to the fully connected layer.

2. Dense Layer: The final Dense layer is a fully connected layer that outputs the predicted stock price for the next day. The loss function used for training is the Mean Squared Error (MSE), a standard loss function for regression tasks in neural networks, particularly when the goal is to minimize the difference between predicted and actual values.
3. Early Stopping: To prevent overfitting and optimize training, an EarlyStopping callback was incorporated into the training process. The EarlyStopping mechanism halts training if there is no improvement in the validation loss after 10 consecutive epochs, restoring the model weights that achieved the lowest validation loss. This technique ensures that the model generalizes well to unseen data while avoiding unnecessary overfitting.

2.3. Training and Evaluation

The RNN model was trained using the training dataset over 100 epochs with a batch size of 32. The Adam optimizer, a widely used adaptive learning rate optimization algorithm, was employed to update the model weights during training. The model's performance was assessed using the following evaluation metrics:

1. Root Mean Squared Error (RMSE): This metric calculates the square root of the average squared differences between predicted and actual values, providing a measure of the model's accuracy by reflecting the typical magnitude of prediction errors.
2. Mean Absolute Percentage Error (MAPE): MAPE measures the average percentage error between the predicted and actual values, offering insight into the model's accuracy in terms of relative error, which is particularly useful for understanding the proportional size of errors in stock price predictions.

The evaluation process involved comparing the model’s predictions against the actual stock prices in the test dataset. Both RMSE and MAPE were calculated to quantify the accuracy and robustness of the model. Performance metrics were used to assess the model’s ability to generalize to unseen data, with lower RMSE and MAPE values indicating better model performance.

3. RESULTS AND DISCUSSION

The evaluation of the RNN model for predicting Taiwan Semiconductor Manufacturing Company (TSMC) stock prices yielded the following performance metrics:

Table 1. Evaluation result table of RNN model

Evaluation Method	Value
RMSE	9.8656
MAPE	5.9017%

These results provide valuable insights into the model's prediction capabilities. The RMSE of 9.87 indicates that, on average, the predicted stock prices deviate by approximately 9.87 units from the actual stock prices. Given that stock prices can fluctuate significantly in volatile markets, this level of error suggests a reasonable level of prediction accuracy. Furthermore, the MAPE value of 5.90% demonstrates that the average percentage error between the predicted and actual prices is below 6%, indicating a relatively low error rate in percentage terms. These metrics collectively imply that the RNN model is capable of capturing the underlying patterns in the stock price data and can generate predictions with a modest degree of accuracy.

Despite the overall positive performance, it is important to examine the results in more detail to fully understand the model's strengths and limitations. The relatively low RMSE suggests that the model performs well in terms of minimizing absolute error, which is crucial in financial forecasting where even small deviations can have significant implications for investment decisions. The RMSE value indicates that the RNN is particularly effective at predicting stock prices in stable or moderately fluctuating periods. However, during periods of high volatility, the model may struggle to maintain the same level of precision. This limitation is common in time series models, where extreme price movements may not be adequately captured by the underlying learning algorithm.

In terms of percentage-based error, the MAPE value of 5.90% reflects a favorable model performance. A MAPE value below 10% is generally considered acceptable in stock price forecasting, and a value below 6% demonstrates that the RNN model is able to maintain a low level of error relative to the actual stock price. This suggests that the model is reliable for generating predictions that are proportional to the actual stock price movements. Nonetheless, while MAPE provides a useful indication of the model’s accuracy, it may be influenced by extreme values or outliers in the data. Therefore, further analysis should focus on how the model performs across different ranges of stock price movements to ensure that it generalizes well to both low and high price environments.

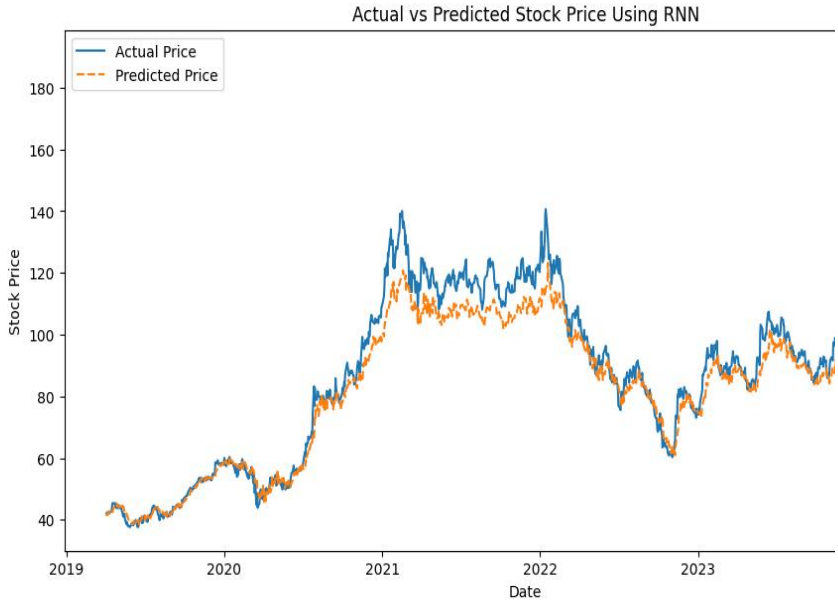

Figure 1. Actual vs Predicted Stock Price using RNN

A key aspect of the analysis involved visualizing the predicted stock prices against the actual prices over the test period. As shown in Figure 1, the RNN model successfully captures the overall trends in the stock price movements. The predicted prices closely follow the actual price trajectory, particularly during periods of steady growth or decline. This alignment between the predicted and actual values indicates that the RNN model is capable of identifying long-term trends and sustaining a consistent predictive performance.

However, the model shows some limitations when predicting stock price behavior during periods of high volatility. As seen in the graph, there are noticeable discrepancies between the predicted and actual prices during sharp upward or downward price movements. These errors could be attributed to the model's inability to fully capture the complex non-linear dynamics that often occur during volatile market conditions. In these instances, the RNN may be limited by its architecture, as the temporal dependencies in the data might not fully account for sudden shifts in market sentiment or external shocks that influence stock prices. Therefore, additional techniques such as incorporating more advanced RNN variants like Long Short-Term Memory (LSTM) networks, or integrating external market factors, could potentially improve model performance during these periods.

The use of early stopping during training proved beneficial in preventing overfitting. By monitoring the validation loss and halting training once no significant improvement was observed, the model was able to retain its generalizability without being overly fine-tuned to the training data. This is particularly important in financial forecasting, where overfitting to historical data could result in poor future predictions due to market changes and unseen events.

In summary, the RNN model demonstrates strong predictive capabilities in forecasting TSMC's stock prices. While the model performs well in capturing general trends, there are limitations in handling high volatility. The relatively low RMSE and MAPE indicate that the model is reliable, but further improvements may be necessary to enhance its robustness in more dynamic market conditions. Future work could explore the integration of additional features or more sophisticated model architectures to better account for market fluctuations and improve accuracy during periods of extreme volatility.

4. CONCLUSION

This research shows that the Recurrent Neural Network (RNN) algorithm can be used effectively to predict stock prices, especially for Taiwan Semiconductor Manufacturing Company (TSMC) stocks. The model evaluation results show that RNN is able to provide accurate predictions with an RMSE value of 9.87 and a MAPE of 5.90%. The RNN model demonstrated the ability to capture the temporal patterns in the stock price data well, resulting in predictions that were close to the actual values. Although the results are already quite good, there is still potential for further improvement with the application of more in-depth preprocessing techniques, hyperparameter tuning, or exploration of other models such as Long Short-Term Memory (LSTM) or Gated Recurrent Units (GRU). Further applications and improvements in this technique may provide more accurate and useful results for future stock market analysis.

REFERENCES

[1] D. D. W. Praveenraj, M. Pandey, and M. Victor, "Time Series Forecasting of Stock Market Volatility Using LSTM Networks," in *2023 4th International Conference on Computation, Automation and Knowledge Management (ICCAKM)*, 2023, pp. 1–8.

[2] P. D. Soni, Y. Tewari, and D. Krishnan, "Machine Learning Approaches in Stock Price Prediction: A Systematic Review," *J. Phys.*, vol. 2161, no. 1, p. 12065, 2022, doi: 10.1088/1742-6596/2161/1/012065.

[3] T. Aamodt, "Predicting Stock Markets with Neural Networks." 2015.

[4] A. Chintam, G. R. R. Kumar, D. Chandramouli, J. A. Rani, and M. A. Praveen, "Accurate stock prices prediction on Grouped Time Series Data using Recurrent Neural Network Variants," 2022, pp. 1–6. doi: 10.1109/ICAECT54875.2022.9807842.

[5] Z. Bao, Q. Wei, Q. Wei, T. Zhou, X. Jiang, and T. Watanabe, "Predicting stock high price using forecast error with recurrent neural network," vol. 6, no. 1, pp. 283–292, 2021, doi: 10.2478/AMNS.2021.2.00009.

[6] A. K. Parida, M. Rout, A. K. Jena, N. Parida, and R. K. Parida, "An Hybrid Stacked LSTM Approach to Forecast the Stock Prices," 2023, pp. 411–415. doi: 10.1109/ocit59427.2023.10430898.

[7] X. Zheng, "Stock Price Prediction Based on CNN-BiLSTM Utilizing Sentiment Analysis and a Two-layer Attention Mechanism," *Adv. Econ. Manag. Polit. Sci.*, 2023, doi: 10.54254/2754-1169/47/20230369.

[8] A. P. da Silva *et al.*, "Performance Evaluation of Recurrent Neural Networks for Short-Term Investment Decision in Stock Market," 2019, pp. 247–261. doi: 10.1007/978-981-15-3852-0_16.

[9] C. Peng, Z. Yin, X. Wei, and A. Zhu, "Stock Price Prediction based on Recurrent Neural Network with Long Short-Term Memory Units," 2019. doi: 10.1109/ICESI.2019.8863005.

[10] C.-C. Kao and C.-Y. Chianglin, "Applying Three Deep Learning Techniques to Predicting Stock Price," 2022, pp. 132–136. doi: 10.1109/IEEM55944.2022.9989878.

[11] C. Duan and W. Ke, "Advanced Stock Price Prediction Using LSTM and Informer Models," *Deleted J.*, vol. 5, no. 1, pp. 141–166, 2024, doi: 10.60087/jaigs.v5i1.183.

[12] "Stock Price Prediction Using Stacked LSTM," *Int. J. Sci. Technol. Eng.*, vol. 11, no. 5, pp. 3578–3584, 2023, doi: 10.22214/ijraset.2023.51757.

[13] H. Harikrishnan and S. Urolagin, "Prediction of Stock Market Prices of Using Recurrent Neural Networkâ€"Long Short-Term Memory," 2021, pp. 359–368. doi: 10.1007/978-981-15-5243-4_33.

[14] I. S. Shcherbyna, "Improving stock price forecasting based on recurrent neural networks," *Nauk. Zap. DerÅ¾avnogo unÃ¬versitetu telekomunÃ¬kacÃ¬j*, vol. 4, no. 2, 2023, doi: 10.31673/2518-7678.2023.021313.

[15] J. V Dave, S. Porwal, U. Jain, G. Chandore, and A. Jain, "Comparing different Machine Learning Algorithms in a stock Market Scenario to check which one has the highest efficiency," 2024, doi: 10.21203/rs.3.rs-4328509/v1.

[16] M. E. Karim and S. Ahmed, "A Deep Learning-Based Approach for Stock Price Prediction Using Bidirectional Gated Recurrent Unit and Bidirectional Long Short Term Memory Model." 2021. doi: 10.1109/GCAT52182.2021.9587895.

[17] N. Naik and B. R. Mohan, "Study of Stock Return Predictions Using Recurrent Neural Networks with LSTM," 2019, pp. 453–459. doi: 10.1007/978-3-030-20257-6_39.

[18] Z. Wang, "Stock price prediction using LSTM neural networks: Techniques and applications," *Appl. Comput. Eng.*, vol. 86, no. 1, pp. 294–300, 2024, doi: 10.54254/2755-2721/86/20241605.

[19] Y. Fu and H. Xiao, "Stock Price Prediction Model based on Dual Attention and TCN," 2022. doi: 10.5121/csit.2022.122007.

[20] X. Zhang, C. Li, and Y. Morimoto, "A Multi-factor Approach for Stock Price Prediction by using Recurrent Neural Networks," vol. 8, no. 1. pp. 9–13, 2019. [Online]. Available: http://www.bncss.org/index.php/bncss/article/download/96/100