# Stock Price Prediction Using Convolutional Neural Networks on Various Time Frames

Jiayao Wang *

Department of Computer Science, Virginia Polytechnic and State University, Blacksburg, VA 24061, United States

* Corresponding Author Email: jiayao@vt.edu

**Abstract.** The stock market is a place where stocks can be transferred, traded, and circulated. It has a history of 400 years and can be used as a channel for companies to raise funds. In this study, two Convolutional Neural Network (CNN) models are developed to forecast stock market prices, catering to distinct investment strategies: a short-term weekly model and a medium-term approximately three-month model. The weekly model utilizes a structure to predict Friday closing prices based on the daily closing prices from the preceding week, intending to capture the weekly trends. In contrast, the medium-term model is designed to comprehend the broader market movements by predicting the closing price on a day nearly twelve weeks ahead, using the past twenty days' prices. Both models incorporate a Conv1D layer followed by batch normalization, pooling, and dense layers, calibrated to their respective temporal frameworks. The study's findings demonstrate the challenges and intricacies of stock price prediction over varying time frames and propose enhancements for future models, including the incorporation of a wider array of market indicators and economic factors.

**Keywords:** Convolutional neural network; stock price prediction; time series data.

## 1. Introduction

Investing in the stock market is a multifaceted decision- making process that encompasses various strategies and timelines [1]. Each investor navigates the market's unpredictable waters with a strategy tailored to meet specific financial goals and risk tolerances [2]. This research is meticulously curated to unveil the potentials of Convolutional Neural Networks (CNN) in forecasting stock market prices, aiding investors in navigating their investment decisions with enhanced precision and confidence [3]. The focal point of this study lies in two predominant investment strategies: weekly-term investments and near three-month term (or 12- week) investments, each wielding its unique approach and time horizon.

The weekly-term investment strategy is structured with a microscopic focus on the weekly stock price movements. Essential attributes such as the closing prices of each week- day are meticulously leveraged, facilitating a detailed analysis that harnesses the weekly fluctuations and trends in the stock market. By focusing on a shorter time horizon, this approach seeks to empower investors with insightful data, allowing them to make informed investment decisions on a weekly basis.

On the flip side, the study also delves into a broader investment strategy, spanning nearly three months or 12 weeks. This strategy unfolds a macroscopic view of the market, absorbing more extended periods of fluctuations and trends. It enables the encapsulation of more comprehensive market movements, providing a broader perspective that is instrumental for investors aiming for a more extended investment horizon.

Central to this research is the application of CNN models, each tailored to resonate with the corresponding investment strategy. CNNs are renowned for their prowess in pattern recognition and have been extensively applied in various domains, including image and sequence data analysis [4, 5]. Their application in this study is aimed at unraveling the hidden patterns and intrinsic correlations in historical stock price data, fostering predictive models that are both robust and insightful. Incorporating the crucial insights provided by the multifaceted approach of artificial intelligence, particularly the prowess of deep learning in automatic feature extraction and prediction, enhances the precision of stock market forecasting, which is central to economic growth and the functionality of effective trading systems [6-8].

The objectives of this research are multifaceted. Primarily, it seeks to investigate the efficacy of CNN models in forecasting stock market prices, examining their predictive accuracy and reliability in the face of the stock market's inherent volatility. Furthermore, the research aims to pro- vide investors with a powerful tool, fortified with accurate predictions and recommendations, ultimately guiding them towards optimized investment decisions across various investment horizons.

## 2. Method

In this research, two distinct Convolutional Neural Networks (CNN) models are meticulously constructed to forecast stock market prices, focusing on two diverse investment strategies: a weekly-term and a nearly three-month term, aligning with approximately bi-monthly or 12-week periods.

### 2.1. Weekly-Term Investment Strategy

#### 2.1.1. Data cleaning

For the weekly-term investment model, historical stock price data were curated with a focus on the daily closing prices. The data was transformed such that each row in the dataset represented a Friday, and included additional attributes detailing the closing prices of the preceding Monday, Tuesday, Wednesday, and Thursday. This transformation aimed to create a structured dataset where the closing price of each Friday was the target variable (Y), while the closing prices of the preceding weekdays acted as the feature set (X), utilized to make the prediction. This approach's intention was to leverage the weekly closing price trends and fluctuations, facilitating the CNN model to capture and learn intrinsic patterns and correlations to predict the subsequent Friday's closing prices.

#### 2.1.2. Model structure

Within the crucible of the weekly-term investment strategy, a CNN model is forged, embodying layers meticulously architected to resonate with the temporal rhythms and nuances of weekly stock market fluctuations [9]. Its structure is demonstrated in Fig 1. The model embarks on its predictive odyssey by initially engaging with a Conv1D layer. This layer, equipped with 32 filters and a kernel size of one, delves into the intricacies of the input sequences, each resonating with the essence of weekly trading attributes.

Model: "sequential"

| Layer (type) | Output Shape | Param # |
|---|---|---|
| conv1d (Conv1D) | (None, 4, 32) | 64 |
| max_pooling1d (MaxPooling1D) | (None, 2, 32) | 0 |
| flatten (Flatten) | (None, 64) | 0 |
| dense (Dense) | (None, 32) | 2080 |
| dense_1 (Dense) | (None, 16) | 528 |
| dense_2 (Dense) | (None, 1) | 17 |

Total params: 2689 (10.50 KB)
Trainable params: 2689 (10.50 KB)
Non-trainable params: 0 (0.00 Byte)

**Fig 1.** Weekly-term model summary (Figure Credits: Original).

Following the Conv1D layer, the model navigates through the realms of normalization and pooling. A Batch- Normalization layer ensures that the model maintains a stable distribution of activations, fostering a conducive environment for the optimization process. Subsequently, a Max-Pooling1D layer with a pool size of two allows the model to distill essential features, navigating through the spatial hierarchies of the input sequences.

As the journey unfolds, the model engages with a Flatten layer, transitioning its understanding from the multidimensional realms to a flattened landscape, preparing the ground for the fully connected layers. The ensuing layers, dense layers with 32 and 16 neurons, respectively, allow the model to weave its learned understanding into predictions, culminating in a final dense layer with one neuron, resonating with the prediction of Friday's closing stock prices.

## 2.2. Three-Month Term Investment Strategy

### 2.2.1. Data cleaning

For the nearly three-month term model, a more extended period was embraced, aiming for a broader and comprehensive understanding of the stock price movements. The data manipulation involved creating new attributes representing the closing prices for each day, from Day 1 to Day 20, embodying almost four trading weeks. Additionally, an attribute named 'Future Friday' was integrated, representing the closing prices 12 weeks post the 'Day 20', acting as the label or the value to predict. An auxiliary attribute 'Future Friday Date' was also incorporated, storing the corresponding dates, facilitating the plotting and visualization process.

### 2.2.2. Model structure

Navigating the broader horizons of an approximate three-month term investment, another CNN model is proposed, echoing with the essence of longer-term market rhythms and vibrations. As shown in Fig 2, this model embarks on its journey through a Conv1D layer with 32 filters and a kernel size of one, fostering an environment conducive to exploring the temporal patterns and nuances across the twenty days of trading data.

| Layer (type) | Output Shape | Param # |
|---|---|---|
| conv1d_1 (Conv1D) | (None, 4, 32) | 64 |
| max_pooling1d_1 (MaxPooling1D) | (None, 2, 32) | 0 |
| flatten_1 (Flatten) | (None, 64) | 0 |
| dense_3 (Dense) | (None, 16) | 1040 |
| dense_4 (Dense) | (None, 1) | 17 |

Total params: 1121 (4.38 KB)
Trainable params: 1121 (4.38 KB)
Non-trainable params: 0 (0.00 Byte)

**Fig 2.** Three-month term model summary (Figure Credits: Original).

Navigating beyond the convolutional layer, the model engages with a MaxPooling1D layer with a pool size of two. This layer fosters a perspective that allows the model to dis- till essential temporal features, reducing the dimensionality while retaining the essence of the longer-term market movements.

Transitioning through the realms of the model, a Flatten layer serves as a bridge, guiding the model's understanding from the multidimensional spaces to a flattened horizon, setting the stage for the fully connected realms. Here, a dense layer with 16 neurons allows the model to craft intricate pathways of understanding, forging connections that resonate with the predictive aspirations of the model. Culminating its journey, the model arrives at a final dense layer with one neuron, embodying the predictions that align with the closing stock prices of a day, twelve weeks into the future.

# 3. Results

## 3.1. Weekly-Term Analysis

### 3.1.1. Comparative analysis of predicted and actual closing prices

The visual comparison between the predicted and actual closing prices for the weekly-term financial analysis is il- lustrated in Fig 3. The depicted graph indicates a no- table congruence between the two data series, signifying the model's effectiveness at capturing the market's trend on a weekly basis. Although there are areas where the predicted values diverge from the actual prices, the general trajectory is closely mirrored. The slight discrepancies be- tween predicted and actual values might be attributed to market volatility or model limitations in capturing sudden price movements. However, the overall predictive accuracy suggests that the model can be a reliable tool for estimating future price trends on a weekly scale.

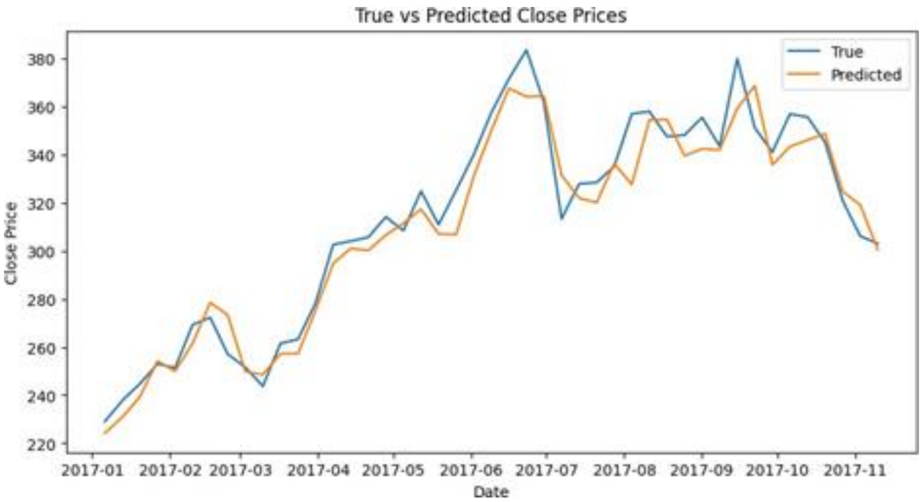


**Fig 3.** Predicted vs. actual closing prices plot (Figure Credits: Original).

### 3.1.2. Influence of weekdays on price predictions

Fig 4 presents the feature importance of weekdays in the predictive model. The uniformity across Monday to Thursday suggests a balanced influence of each day on the model's predictions. The slight variations in importance scores are likely inconsequential, indicating that the day of the week may not be a strong determinant in the model's forecasting ability for this particular dataset. This uniform distribution hints at the market's efficiency, reflecting that no single weekday possesses a significantly higher predictive power over another within the scope of this model.

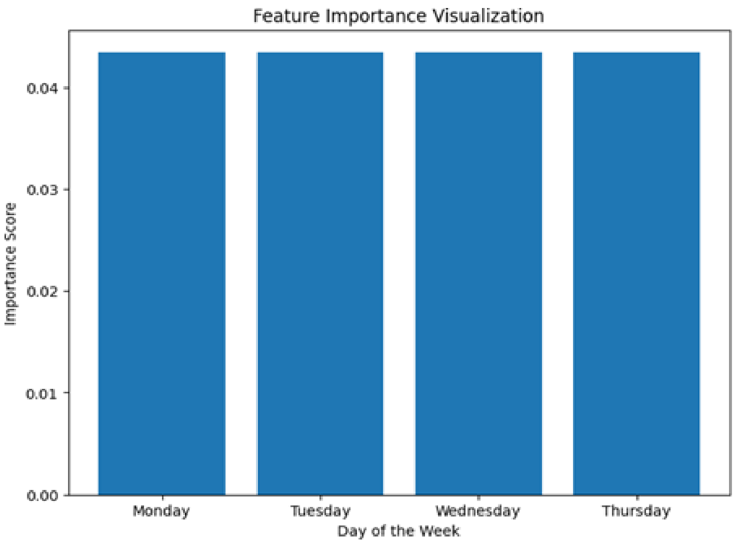


**Fig 4.** Feature importance visualization (Figure Credits: Original).

### 3.1.3. Comparison of actual vs predicted weekly volatility

Fig 5 illustrates the comparison between the actual and forecasted weekly volatility. The plotted lines reflect the model's capacity to trace the general trend of the market's fluctuations, with the predicted volatility often mirroring the direction of the true volatility. Notably, there are periods where the predicted values diverge from the true ones, highlighting potential areas for model refinement. The model's adeptness in tracking the peaks and troughs of the market volatility suggests its utility in risk assessment and strategic decision-making processes.

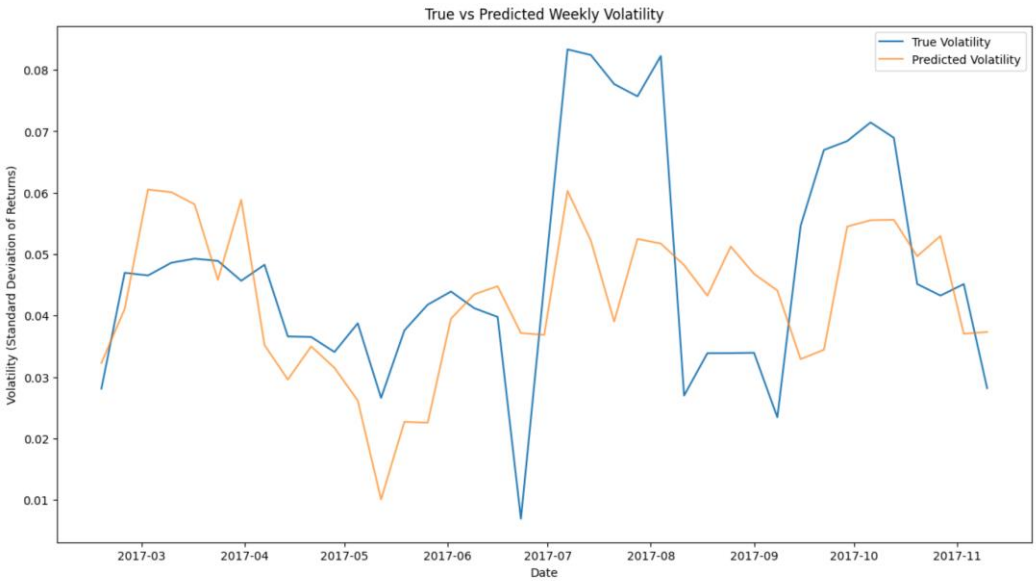


**Fig 5.** True vs predicted weekly volatility (Figure Credits: Original).

### 3.1.4. Investment return simulation

In this simulation, an initial investment of $100 was sub- ject to two different investment strategies over a series of weeks. The first strategy was based on actual weekly closing prices (True Prices), and the second strategy was based on predicted weekly closing prices (Predicted Prices).

The simulation yielded two different growth trajectories for the initial investment. When the strategy was executed using the true closing prices, the investment's value increased to $271.01 by the end of the period. This represents an ideal scenario with perfect foresight into weekly price movements, which, while not achievable in practice, serves as a benchmark.

In contrast, when the investment strategy was based on predicted prices, the investment grew to $212.93, demonstrating that the predictive model was able to capture the general trend of the price movements, but with less precision than the actual prices. The difference in the final investment values reflects the challenges and uncertainties associated with financial forecasts.

Despite the lower final value compared to the true price strategy, the strategy based on predicted prices did realize a substantial return, underscoring the utility of predictive models in investment strategies, albeit with the acceptance of a margin of error.

The comparative results of this simulation illustrate the potential and limitations of using predictive models in financial decision-making. While they cannot replicate the success of an approach with perfect information, they still provide valuable insights that can lead to above-baseline returns on investment.

### 3.1.5. Confidence interval of predicted values

The analysis of predicted financial metrics often incorporates a measure of uncertainty, typically expressed as a confidence interval. The confidence interval offers a range within which the true values are expected to fall, with a certain level of confidence, usually 95.

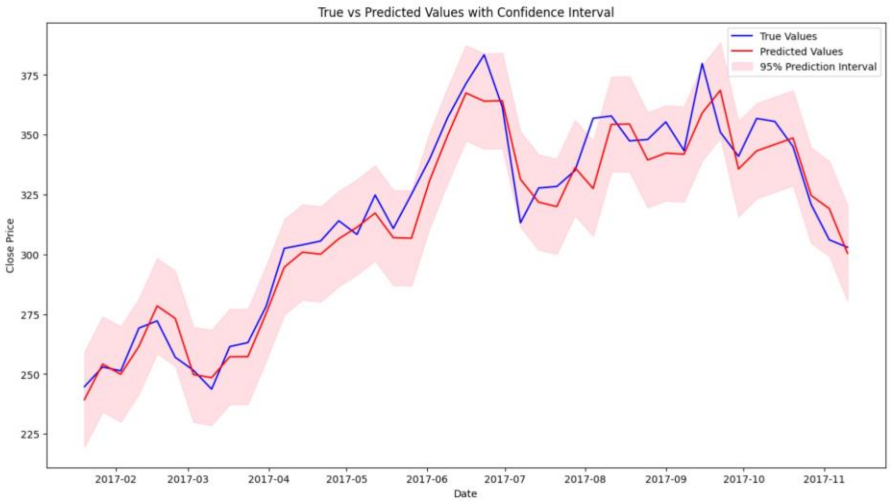


**Fig 6.** True vs predicted values with confidence interval (Figure Credits: Original).

The chart depicted in Fig 6 demonstrates the true values of an unidentified financial metric in comparison to the predicted values over the same period, with the shaded area representing the 95.

This divergence underscores the presence of unpredicted volatility and the limitations inherent in the predictive model used. The instances where the true values exceed the bounds of the confidence interval may be attributed to sudden market movements or unforeseen events impacting the financial metric.

The ability of the model to encompass the true values within the confidence interval for the majority of the time suggests a generally reliable predictive performance. However, investors and analysts must consider the potential for outlier events and the associated risk when making decisions based on these predictions.

### 3.2. Twelve weeks-term analysis

The Twelve Weeks-Term Analysis chart (Fig 7) offers a visual comparison between the true values and the predictions made by the model for a 20-day term. The depicted values illustrate the model's limited ability to predict over the specified term, with the 'Day20' line representing the end of the predictive period. The actual values (blue line) fluctuate significantly compared to both the predicted (red line) and 'Day20' (green line) values, which indicates considerable discrepancies between expected and real out- comes. These results suggest that the model may require further refinement or additional data to enhance its predictive accuracy for this specific time horizon.

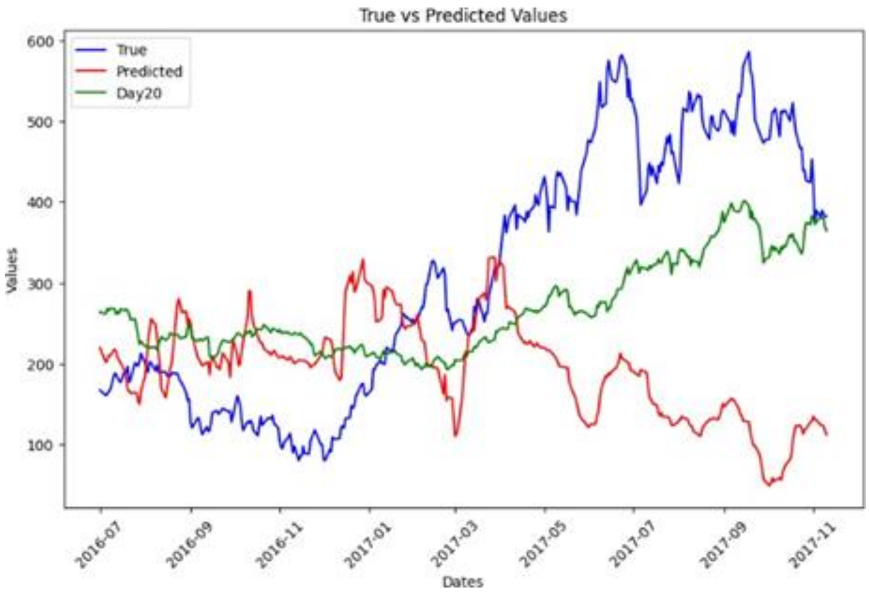


**Fig 7.** Comparison of true vs predicted values over a twelve weeks-term (Figure Credits: Original).

## 4. Discussion

The results obtained from the predictive model, particularly in the Twelve Weeks-Term Analysis, indicate discrepancies between the true values and predicted outcomes. Several factors could contribute to these variances. Firstly, the model's predictive power for the 20-day term appears to be limited, as evidenced by the lack of alignment between the predicted and actual values. It is hypothesized that this might be due to the inherent volatility in the dataset that was not fully captured by the model's parameters. More- over, external factors not included in the dataset, such as market sentiment or unforeseen economic events, could also account for the differences observed [10].

Given the results, it is advisable to integrate additional variables that may better account for market dynamics. Incorporating factors such as volume trends, news sentiment analysis, or even macroeconomic indicators could potentially enhance the model's predictive accuracy. Furthermore, exploring alternative modeling techniques or more complex neural network like long short-term memory (LSTM) architectures could yield improvements in the 20-day forecast horizon.

One of the primary limitations of this study is the model's reliance on price data alone. The financial markets are influenced by a complex interplay of factors, and the model's current scope may not adequately capture these elements. Another limitation is the model's short-term prediction horizon, which does not account for long-term trends or cyclical patterns. To address these issues, future research could aim to incorporate a broader set of predictive variables and perhaps adopt a hybrid modeling approach that combines both quantitative and qualitative data. Additionally, extending the prediction horizon and testing the model across various market conditions could provide more insight into its robustness and reliability.

## 5. Conclusion

This research evaluates two Convolutional Neural Networks (CNN) designed for short-term (weekly) and medium-term (approximately three-month) stock price predictions. The weekly model uses the closing prices from Monday to Thursday to predict Friday's closing prices. It starts with a Conv1D layer, followed by batch normalization and max pooling, and concludes with densely connected layers, including a final single neuron output layer predicting the upcoming Friday's closing price.

For the nearly three-month term model, attributes are created to represent daily closing prices over four weeks, and a 'Future Friday' label representing the price 12 weeks later. This model also begins with a Conv1D layer and includes max pooling and flattening before concluding with densely connected layers.

The results demonstrate that the weekly model closely mirrors actual price trends, with minor discrepancies attributed to market volatility or model limitations. It effectively predicts weekly price movements and volatility, proving useful for risk assessment and strategic decision-making. An investment return simulation using the weekly model shows substantial returns, though not as high as an ideal scenario with perfect market foresight.

For the three-month model, there are significant differences between predicted and actual values, suggesting the need for refinement. Confidence intervals show the model's general reliability, but also reveal unpredicted volatility and inherent limitations.

The study concludes that while the models offer valuable insights for investment strategies, there is room for improvement. Future work could integrate more variables to account for market dynamics and adopt a hybrid modeling approach to enhance predictive accuracy. The research suggests expanding the prediction horizon and testing across various market conditions to better understand the models' robustness.

## References

[1] Lu, Wenjie, Jiazheng Li, Jingyang Wang, and Lele Qin. A CNN-BiLSTM-AM method for stock price prediction. Neural Computing and Applications, 2021, 33: 4741-4753.

[2] N. Samuel. Changing market conditions impact your asset allocation strategy. Picture Perfect Portfolios, 2023. URL: https://pictureperfectportfolios. Com / how - changing - market - conditions - impact-your-asset-allocation-strategy/. Last access time: 2023/11/09.

[3] Gu, Jiuxiang, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu et al. Recent advances in convolutional neural networks. Pattern recognition, 2018, 77: 354-377.

[4] Rawat, Waseem, and Zenghui Wang. Deep convolutional neural networks for image classification: A comprehensive review. Neural computation, 2017, 29(9): 2352-2449.

[5] Al-Saffar, Ahmed Ali Mohammed, Hai Tao, and Mohammed Ahmed Talab. Review of deep convolution neural network in image classification. In 2017 International conference on radar, antenna, microwave, electronics, and telecommunications, 2017: 26-31.

[6] Hoseinzade, Ehsan, and Saman Haratizadeh. CNNpred: CNN-based stock market prediction using a diverse set of variables. Expert Systems with Applications, 2019, 129: 273-285.

[7] Hiransha, Ma, E. Ab Gopalakrishnan, Vijay Krishna Menon, and K. P. Soman. NSE stock market prediction using deep-learning models. Procedia computer science, 2018, 132: 1351-1362.

[8] Selvin, Sreelekshmy, R. Vinayakumar, E. A. Gopalakrishnan, Vijay Krishna Menon, and K. P. Soman. Stock price prediction using LSTM, RNN and CNN-sliding window model. In 2017 international conference on advances in computing, communications and informatics, 2017: 1643-1647.

[9] Mehtab, Sidra, and Jaydip Sen. Stock price prediction using CNN and LSTM-based deep learning models. In 2020 International Conference on Decision Aid Sciences and Application, 2020: 447-453.

[10] Rasheed, Jawad, Akhtar Jamil, Alaa Ali Hameed, Muhammad Ilyas, Adem Özyavaş, and Naim Ajlouni. Improving stock prediction accuracy using cnn and lstm. In 2020 International Conference on Data Analytics for Business and Industry: Way Towards a Sustainable Economy, 2020: 1-5.