AI-Driven Stock Return Prediction: Evaluating CNN, LSTM, and RF for Nvidia

Zonghan Wu[a]
Bill Hogarth Secondary School, Markham, Canada

Keywords: Stock, Machine Learning, Artificial Intelligence.

Abstract: This paper presents a new approach for estimating Nvidia stock returns using advanced learning algorithms, including Convolutional Neural Network (CNN), Longg Short Term Memory (LSTM), and Random Forest (RF). The system methodology focuses on identifying complex market dynamics by analyzing daily stock returns. Features are preprocessed through normalization to stabilize variance. The CNN architecture involves three 1-D convolutional layers with 64, 128, and 256 filters to scan temporal patterns, followed by two LSTM layers with 50 neurons each to capture long-term dependencies. Random Forest with 100 trees balances computational complexity and predictive performance. Models are trained on 80% of the data, with 20% reserved for testing. Evaluation results indicate that the LSTM model outperforms CNN and Random Forest based on RMSE and MAE metrics. However, the models do not account for external factors like news events and economic indicators, limiting predictability. This study demonstrates the effectiveness of LSTM in predicting stock returns and lays the groundwork for future enhancements in AI-based financial models, with potential applications in algorithmic trading and risk management.

1 INTRODUCTION

Stocks are normally considered one of the most prominent methods for investment. They are normally referred to as a form of ownership in a corporation and embody claims on part of its assets and earnings. Their importance cannot be measured only by individual benefits, which they make possible, because they are quite important for any nation or person to experience economic growth and financial stability. Hence, the prediction of stock prices becomes very important not only to the investors but also to the financial analysts, the makers of policies, and strategists of economies. However, traditional methods for the prediction of stock prices have depended on fundamental and technical analysis making use of time series analysis, regression models, and statistical indicators. These conventional methods always turn out to be inadequate in the presence of high volatility, non-linearity, and complex patterns of the stock market data, hence yielding suboptimal accuracy in prediction. In this case, more advanced models with superior performance should be considered.

During the last few years, Artificial Intelligence (AI) has moved to such a great extent that its influence is not only deeply attached to health and robotics but also to natural language processing and financial forecasting (Szolovits, 1988; Holmes, 2004; Miller, 2018; Roll, 2016). Hence, it becomes especially applicable to the prediction of stock prices, since AI can process large amounts of data and capture complex patterns to deal with nonlinear relations and changes that are dynamic in the market. Of the different methodologies associated with artificial intelligence, those based on machine learning have had enormous successes, particularly on those applying deep learning algorithms. For instance, in 2018, Fischer et al. applied Long Short Term Memory (LSTM) networks in predicting the directional movements of S&P 500 stocks, like moving averages and the relative strength index (Fischer, 2018). This improved a great deal the accuracy of the prediction, thus verifying that there were benefits expected to be derived when technical indicators were combined with AI models. Similarly, in 2017, Bao et al. reviewed the efficacy of LSTM versus support vector machines and random forests in

[a] https://orcid.org/0009-0007-4518-5474

Wu and Z.
AI-Driven Stock Return Prediction: Evaluating CNN, LSTM, and RF for Nvidia.
DOI: 10.5220/0013487800004619
In *Proceedings of the 2nd International Conference on Data Analysis and Machine Learning (DAML 2024)*, pages 80-84
ISBN: 978-989-758-754-2

stock price prediction (Bao, 2018). The conclusion was that LSTM is better because of improvement in capturing the ability of complex time-series data. These examples provide evidence of the utility that AI models could add to the prediction of stock prices in the future.

Not only these particular studies but the general financial community has also realized the potential of AI. Many algorithms have been used, including convolutional neural networks and reinforcement learning, and also hybrid models mixing these different techniques in order to improve accuracy in prediction. For example, Kumar et al. used a hybrid model mixing LSTM with convolutional neural networks in stock price prediction (Krauss, 2017). The approach was found to capture the time dependencies and local features very well. Coupled with increasing accessibility to the actual financial data, these developments have begun to yield more sophisticated and reliable models for prediction, which financial institutions and individual investors are increasingly adopting.

The paper aims to use the advanced AI techniques in predicting the NVIDIA stock price in view of some encouraging results on the application of AI models towards stock price forecasting. The long short-term memory and convolutional neural network, and hybrid approaches called random forest will compare multiple models involved in view of their efficiency for predicting stock prices. Models will also be judged based on mean squared error, mean absolute error, and R-squared to ensure that models being developed get a full judgment about predictive capabilities. More precisely, these AI-driven methodologies can make this study part of the emerging discourse regarding betterment in stock price prediction and provide valuable insights that help investors, analysts, and researchers.

2 METHOD

2.1 Dataset Preparation

This research paper uses data that was gathered from Yahoo Finance (Yahoo, 2024), covering the period from 1999 to 2024. The variable used for this research paper was the daily return of NVIDIA stock computed as a percentage change in closing price over one day. This variable is of importance because the percentage return actually takes into consideration the fluctuation in the price over a day; hence, it provides a good snapshot of market volatility and investors' sentiment compared to the unprocessed closing price.

The original dataset included daily financial indicators: Open, High, Low, Close, Adjusted Close, and volume metrics. The data went through a comprehensive preprocessing before feeding into the machine learning model. Missing values were first identified and then corrected, either by imputation or exclusion, up to this point in time and date for integrity. Robust statistical methods were also employed for identifying and treating any identified outliers so that they do not overly influence the learning of the model.

Therefore, the datasets were split into 80% for training and 20% for testing. This way, the models can have enough samples to learn from while reserving some samples for testing the performance of the trained models on totally unseen sets of samples. Because the present dataset is a time series dataset, this split will be carried out sequentially in order to retain its temporal structure.

In addition, normalization was applied on the dataset by min-max scaling. It rescales all values of each feature into a range from 0 to 1. For models like Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) (Li, 2021; Graves, 2012), it may be important because differences in several features' scales act as a barrier to convergence and reduce model performance.

2.2 Predictions Using Machine Learning Models

This paper considers three different machine learning models for prediction: Convolutional Neural Network, Long Short-Term Memory, and Random Forest. These were chosen based on their different performances in trying to overcome problems related to time-series data in stock returns. The actual implementation of the CNN and LSTM has been done using TensorFlow, while the Random Forest model has been used with Scikit-learn. Concerning the assessment of the predictive performance, the calculations for Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) have been put forth with respect to all three models in an informative way.

2.2.1 Convolutional Neural Network

This study employs 1D CNN since it can capture the temporal structure of sequential data and thus fits well for stock return forecasting. It's very different from other CNNs, such as 2D; it is dedicated to image data. The 1D CNN process of information in time series is made on the time axis that could easily spot local patterns and trends of the time series.

It was an architectural design consisting of input, huge convolutional layers applied with the Rectified Linear Unit (ReLU) activation function, and max-pooling layers for extraction of features, followed by a dense layer where the prediction is done. The CNN model consisted of three 1D convolutional layers with the configuration 64, 128, and 256 filters, respectively. Each layer was followed by a kernel size of 3 to try to catch up with as many short-term dependencies of the data as possible. Finally, an output dense layer was defined in which one neuron was initialized for predicting the daily return. It will be compiled using the Adam optimizer and mean squared error loss function.

2.2.2 Long Short-Term Memory

Long Short-Term Memory networks are a special type of RNN able to learn long-term dependencies in sequences. Which may prove very efficient since financial time series are usually strongly impacted by historical trends. This study uses the model based on LSTM with the ability to learn such complex dependencies in daily return data in an effort to arrive at better predictive performance. The LSTM architecture was composed of two stacked LSTM layers with 50 neurons each, followed by a dense layer with 100 neurons and the ReLU activation function. This brought more capability of high-order representation learning to the model. The final prediction was done using an output single neuron optimized through an Adam optimizer with a mean squared error loss function.

2.2.3 Random Forests (RF)

The random forest is a technique for ensemble learning in which outputs from multiple decision trees are combined, and the average of their output gives the conclusion for a prediction. This approach is much stronger and highly suitable as it deals with both linear and nonlinear relationships. A possible reason why Random Forest can do a good job may be that through averaging many trees, where each tree has been fitted to different subsets of data, it avoids overfitting. This paper utilized a Random Forest model with 100 decision trees, a reasonable balance between computational efficiency and predictive accuracy. Each tree in the forest is trained on part of the available data, whereas the overall estimation is derived from averaging over all tree outcomes.

The present study will leverage deep learning methodologies, especially CNNs and LSTMs, apart from traditional machine learning techniques such as RF. Those above-mentioned variate methods facilitate stock returns in NVIDIA to be analyzed comprehensively, reflecting not only short-term fluctuation but also long-term trends.

3 RESULTS AND DISCUSSION

3.1 Overview of Experimental Results

This study evaluated the performance of three models—CNN, LSTM, and RF—in predicting daily returns. The key evaluation metrics used were RMSE and MAE. The results indicated that the LSTM model outperformed the other two models in both metrics.

Table 1: The performance comparison among different models.

Model Name	RMSE	MAE
LSTM	0.0337	0.0247
CNN	0.0341	0.0249
RF	0.0342	0.0250

Concretely, the value for RMSE was equal to 0.0337 and MAE equalled 0.0247 for the LSTM model shown in Table 1, which is quite good in terms of daily stock prediction. On the other hand, CNN recorded an RMSE of 0.0341 and MAE of 0.0249, showing a little higher rate of error compared to LSTM. The Random Forest model had the highest error rates, with an RMSE of 0.0342 and an MAE of 0.0250, indicating limitations in capturing the complex patterns in daily returns.

3.2 Data Interpretation

The experimental evidence suggests that an LSTM network really suits the network for the time series forecasting problem, particularly in the domain of stock return prediction tasks. The actual idea lying behind this good performance works around built-in memory cells that can hold long-term dependences and use them in the data themselves. Long-lived memory is very important for financial time series and mainly suggests the effect of past data in deciding future prices should be large. This is because, for the inherent task, LSTM models manage best and recall the most relevant information with greatest efficacy across time. Although CNNs really excel in local pattern detection, such as with images, they performed rather badly in time series predictions. That is understandable, given the nature of a CNN; generally, it doesn't have the same power that LSTMs do because of their inability to handle long-range dependencies within sequential data. CNNs really perform well when capturing relationships among spatial dimensions but may lag concerning their ability with temporal sequences contained within financial data. Though powerful in handling non-linear relationships, random forests came out as the least effective in the study. The main drawback relates to time-series data, as it can't catch the sequential nature of input. Unlike LSTM and CNN, designed to work with ordered data, a priori, random forests consider each observation independently, which might lead to a loss of the temporal information that could well be critical in making good predictions in financial markets.

3.3 Discussion on Model Performance

These results, therefore, support the current literature only with regard to the strengths and weaknesses within the different models of machine learning applied to time series prediction. The superior performance reached through the LSTM network owes its origin to its specialized design, which enables it to keep information about the past sequences and utilize it better compared to CNN and Random Forest. This turns out to be apt in the case of a stock market forecast since it usually depends on the past historical price patterns, which influence the future or subsequent movements.

While useful insights can be captured by CNNs, it is clear that without the full integration of time-based information, their inclusion has limited benefits compared to an LSTM network. Such inability was further reflected in the more analogous error rates seen in the predictions postulated by the CNN model. Notwithstanding this, CNNs may still prove useful in hybrid models where their focus on different aspects of the data can complement other techniques.

Having seen how this Random Forest model underperformed, it just goes to show the kind of difficulties one gets with using classical machine learning algorithms on time-series data without preprocessing. While Random Forest works wonders in an environment that has complex nonlinear relationships among its many variables, their failure to consider the fact that data points come into clear order makes them unfit for application when the order matters most, such as stock returns.

Although the performance of the LSTM model was satisfactory, some issues led to specific shortcomings of this study. First and foremost, exogenous variables are lacking in the model: economic indicators, news events, and geopolitical events are things that should have particularly affected stock prices. A road furthered by the research would be the incorporation of such exogenous factors into the model for enhanced performance. Besides, the models have been developed on historical data alone, assuming that past trends will continue into the future. However, financial markets are several times swayed by surprising events, and probably future studies can look at models that can capture such variabilities in a better way.

Where an accent on the daily returns is informative-sometimes too little can adequately account for finer fluctuations. Having said this, the next step in the study has to be the establishment of just how effective these models are by using data of higher frequencies, such as hourly returns, showing whether even finer and timelier predictions can be used. Ultimately, attempting to address some of these problems and to find new avenues of research will result in more universal predictive models of financial markets.

4 CONCLUSIONS

This study has utilized advanced machine learning techniques, specifically Convolutional Neural Networks, Long Short-Term Memory networks, and Random Forests, in order to forecast the daily returns of the NVIDIA stock. Using the historical data on stock returns, the study works with an aim to detect some sophisticated patterns in the market and check the predictive power of the considered models. The results of the experiments showed that the LSTM model outperformed both the CNN and RF models in terms of accuracy. Accuracy was measured with RMSE and MAE.

The results of the empirical work developed herein demonstrate the aptness of LSTM for financial forecasting, especially in capturing sequential dependencies embedded in time series. The models being proposed could be used for extending the scope of decision-making activities in algorithmic trading and risk management. However, none of the presented models integrates news events and economic indicators so far; such factors would seriously affect their performance. Much research is required in the future to concentrate on their integration, which would lend greater efficiency and robustness to the models.

REFERENCES

Bao, Y., Li, J., & Zhu, Y. 2018. Mammary analog secretory carcinoma with ETV6 rearrangement arising in the conjunctiva and eyelid. The American Journal of Dermatopathology, 40(7), 531-535.

Fischer, T., & Krauss, C. 2018. Deep learning with long short-term memory networks for financial market predictions. European journal of operational research, 270(2), 654-669.

Graves, A., & Graves, A. 2012. Long short-term memory. Supervised sequence labelling with recurrent neural networks, 37-45.

Holmes, J., Sacchi, L., & Bellazzi, R. 2004. Artificial intelligence in medicine. Ann R Coll Surg Engl, 86, 334-8.

Krauss, C., Do, X. A., & Huck, N. 2017. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. European Journal of Operational Research, 259(2), 689-702.

Li, Z., Liu, F., Yang, W., Peng, S., & Zhou, J. 2021. A survey of convolutional neural networks: analysis, applications, and prospects. IEEE transactions on neural networks and learning systems, 33(12), 6999-7019.

Miller, D. D., & Brown, E. W. 2018. Artificial intelligence in medical practice: the question to the answer?. The American journal of medicine, 131(2), 129-133.

Roll, I., & Wylie, R. 2016. Evolution and revolution in artificial intelligence in education. International journal of artificial intelligence in education, 26, 582-599.

Szolovits, P., Patil, R. S., & Schwartz, W. B. 1988. Artificial intelligence in medical diagnosis. Annals of internal medicine, 108(1), 80-87.

Yahoo. 2024. NVIDIA Corporation (NVDA). Retrieved from https://ca.finance.yahoo.com/quote/NVDA/history/?period1=916963200&period2=1725062400&interval=1d&filter=history&frequency=1d&includeAdjustedClose=true