International Journal of Modern Developments in Engineering and Science
Volume 5, Issue 5, May 2026
www.ijmdes.com | E-ISSN: 2583-3138 | RESAIM Publishing | www.resaim.com

Stock Price Prediction Using a CNN-LSTM Based Deep Learning Approach

H. Karthikeyan[1], S. Prabhu Rajan[2*]
[1]Department of Department of Networking and Communications, SRM Institute of Science and Technology, Chennai, India
[2]Department of Computational Technology, SRM Institute of Science and Technology, Chennai, India

Abstract*: Early and precise stock price forecasting has become a crucial aspect of financial risk reduction and a better approach towards investing in the past few years as a result of the volatility in financial markets. The conventional analytical methods are based on manual interpretation and statistical measures, which can be incapable of quantifying nonlinear relationships existing in the market data. To overcome these drawbacks, we have come up with a hybrid deep learning model of stock price prediction, which is a hybrid of Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks. Technical indicators and subjective estimates are usually useful in the stock price prediction and may be subjective as well as lengthy across various analysts. The data has stock data (Open, High, Low, Close and Volume) of some of the most successful technology companies such as AAPL, AMD, IBM, GOOGL, AMZN, NVDA, EBAY and CSCO on a daily basis. The data preprocessing will involve data cleaning, 20 days moving average (MA20) and Min-Max normalization to give a stabilized training behaviour. These representations are then inputted to LSTM layers which are used to model long term dependencies and sequence relations of price changes of stocks. This stratification ensures that reuse of features is also efficient not to mention that the valuable time information is saved. The network is computationally efficient and can be trained without the use of a lot of hardware hence can be applied in practice. This model was trained and tested on 2019-2024 data related to the stock market. Experimental outcomes show that it has better predictive performance than compared to traditional standalone models with reduced Mean Squared Error (MSE). In order to increase usability, an interactive analysis interface was created with the help of Plotly that allows visualizing the comparison between the predicted and the actual stock prices. The proposed framework offers the solution of financial time-series forecasting based on a scalable and efficient approach and can help investors and analysts to make well-informed decisions. The system is flexible to be incorporated in web-based financial analysis systems and it can facilitate on-the-fly market analysis.

***Keywords*: Deep Learning, CNNLSTM, Stock Price Prediction, Financial Time-Series Forecasting, Technical Indicators, Mean Squared Error, Machine Learning in Finance, Data Driven Investment Analysis.**

1. Introduction

Stock market forecasting has proved to be one of the most challenging and powerful problems in the sphere of financial analytics. Equity markets involve millions of investors in a year and the price movement is highly unpredictable due to volatility, variability of the economy and due to a high percentage of information flow. Any error committed during prediction can result in massive losses of money. The technological stocks are especially dependent on the market sentiment and some events occurring in the world economy, which makes the forecasting reliability of the technology stocks even more vital than in other industries.

It has been demonstrated that predictive modeling can bring a huge improvement in portfolio performance when good predictive models are applied. The good forecasts, however, might require to be done by skilled analysts, by complex instruments of analysis and by constant observation of the market which might not be easily accessible to small investors. In the vast majority of cases, individuals are controlled by mere technical indications or guesses and, therefore, take inconsistent and risky decisions.

Traditional statistical models, such as moving averages and autoregressive, are limited to the ability to model nonlinear relationships and long-run time relationships present on financial time-series data. Although the predictive skills have improved, currently, the machine learning methods cannot usually be used to learn both the short-term and long-term dynamics simultaneously. Furthermore, lone deep learning models either can disregard local finer-grained variations, or can be incapable of preserving sequential dependencies.

Most of the existing financial forecast systems are single model based systems and they are usually not illustrated by demonstration and useful deployment interfaces. This makes them less useful in actual life. Besides, financial time-series data is highly sensitive to noise, prices spikes and fluctuating volumes of trading that also complicates the generalization of the model.

To address these problems, the research paper shall propose a hybrid CNNLSTM model to forecast share prices. The fact as to why the two models are to be combined is because it is complementary to their benefits. The CNN component is able to compress local trend patterns, short-term structural information out of sequential input data but the LSTM one can compress long-term dependence in time and sequential behaviour. This unified structure allows the reuse of features efficiently and allows the training to be stable, and allows the retention of critical information over time steps.

*Corresponding author: prabhurajan63@gmail.com

The model is trained using the historical stock data of such big technology companies as AAPL, AMD, IBM, GOOGL, AMZN, NVDA, EBAY, and CSCO, between 2019 and 2024. The data set contains the characteristics of trading on a daily basis, which are Open, High, Low, Close, and Volume. Evaluation of experiments is superior in predictive performance over the traditional method, in terms of Mean Squared Error (MSE). The suggested system offers a computationally efficient and scalable system that can be incorporated into financial analytics tools in order to facilitate informed decision-making when it comes to investing**.**

2. Literature Review

During the last ten years, a lot of research has been carried out in the field of stock market forecasting owing to the growing need of smart financial forecasting systems. The original study was predominantly founded on statistical techniques such as Moving Averages, ARIMA models and linear regression strategies. They may not have been capable of being able to attain the nonlinear relationship and sudden correlations and transformations in the market of the world economy events and investor sentiment, though these classical models provided the fundamental insight.

To address these lacks, scientists began applying the procedures of Machine Learning (ML) such as Support Vector Machines (SVM), K-Nearest Neighbors (KNN), and Random forests. These methods were more predictive successful in learning patterns that were used on past financial records. Nevertheless, they usually demanded a high degree of feature engineering, as well as human selection of technical indicators. These models were largely unsuccessful when it comes to adjusting to different market conditions and different types of stocks because the credibility of handcrafted features was an important parameter.

Due to the growth of computational resources, Deep Learning solutions became a breakthrough paradigm of financial time- Series forecasting. Conventional application of Convolutional Neural Networks (CNNs) in image processing was also modified to the financial data based on time-series sequences as structured data. CNN-based models showed excellent performance in identification of the local variations and short term trends in prices. However, standalone The CNN models could not capture long-term temporal correlation of stock market data.

Recurrent Neural Networks (RNNs) previously called the Long Short-Term Memory (LSTM) networks have been proposed to resolve sequential (or time-related) learning problems. LMST networks were rather successful in solving the problem of the vanishing gradient and they showed enormous ability in learning long term dependencies. Some of the studies have shown that the forecasting accuracy of LSTM-based architecture is improved. In some cases, though, the LSTM-only based models failed to trace the finer-time fluctuations (they being sequential processors).

The recent studies were conducted on the hybrid deep learning networks which combine CNN and LSTM networks. Features in this type of architecture are gained by the CNN layers and sequential dependency modelling is implemented by the LSTM layers. Mahesh (2009, p. 12) indicates that this type of hybrid techniques have proved to be more predictive measures as well as more stable than some single-model implementations. All these have evolved, but majority of the systems proposed are on the test stage and are not provided with viable interfaces on which to analyse the financial information in real time.

The other issue of high volatility and noise in the stock market data is another associated issue that is a major challenge in financial forecasting. The abrupt variation in price can be extremely sensitive to the market and unfair trading level and market mood that may result in the overfitting of models and inefficient generalization. Despite the mechanisms that usually are applied to enhance the level of robustness, i.e., feature scaling, regularization, and dropout that are used, the opportunity to create the system that will not only allow the improved predictive performance but the number of computational costs that can be plausible is a subject of research.

The fact that it has been considered that the available financial forecasting systems have been significantly improved due to the implementation of the deep learning demonstrates that there are no connections between the research models that are highly accurate and the practical application of the available systems. The gap area in this case results into the creation of a hybrid CNNLSTM model that works excellently in features extraction, sequential representation and performance. The proposed solution will remove these limitations by offering a fair balance solution that will be capable of application to the real world financial decision-support systems.

3. Existing System

Technical analysis and manual use of charts, technical indicators and judgment in traditional financial analysis constitutes an important part of the stock price prediction process. Moving averages, Relative Strength Index (RSI), and MACD are some of the tools used by traders and analysts to make estimates of the possible price changes. These approaches are popular, however, they are also subjective and reliant on the experience and intuition of the analyst in an immense sense. The decisions reached by two financial analysts on the stock data of the same stock, might not be identical at all, and therefore, the investment decisions will not be congruent.

The other major weakness of the traditional forecasting method lies in the fact that it entails the continuous observing and interpreting of the market indicators. Small scale investors lack access to advanced tools of reporting analysis or professional advice. This has contributed to coming up with many speculations or simplified decisions that are not representative of the dynamic market dynamics. This places the company in the predicament of making losses particularly in volatile markets such as technology shares.

Earlier forms of computation attempted to make predictions on a statistical basis using ARIMA and linear regression that were automated. Although these methods provided systematic prediction models, they were limited to nonlinear association and rapid alteration of the market. This was followed by shallow machine learning models which included Support Vector Machines (SVM), K-Nearest Neighbors (KNN) and Decision Trees. Although these methods enhanced the pattern recognition abilities, they needed a great deal of feature engineering and selection of pertinent technical indicators manually.

One of the major problems of these traditional and shallow learning systems is that they use handcrafted features. Researchers will have to clearly define the market signals, trends or indicators that the model will take into consideration. It is typical that this type of extraction process used in the various features of the manual does not generalize to different market conditions, particularly during unforeseen event(s) like an economic recession or a price change spurt. The resultant effect of these systems is that they can make incorrect predictions or have large error rates.

The recent advances have presented individual deep learning models like Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). Although CNN-based models can extract short-term dynamics in ordered sequences of inputs, they might not be able to reason long-term dynamics in financial time-series data. It is also akin that when used as such, RNN-based approaches may prove inadequate in portraying fine-grained local variations because they are likely to be categorized as high-level features during sequence modelling.

Another weakness that has been observed in most of the systems is lack of viable deployment frameworks. The laboratory still retains some of the successful research models and they have not been integrated into scalable and easy to use financial analysis systems. Also, financial data is also reported to be very noisy and volatile and most of the existing models do not adequately address the overfitting and generalization issue of various companies.

In general, the existing stock price prediction methods have shortcomings such as the use of manual feature selection, inability to capture nonlinear temporal patterns, prone to volatility-related errors, and unavailability in the real world. Such weaknesses underscore the necessity of a more combined and powerful forecasting system with the ability to learn both long-term and short-term relationships and to be applicable and economical to compute at the same time.

4. Proposed System

The proposed system will provide an automated and data-driven approach of the stock price prediction through the assistance of the hybrid CNN-LSTM model. Unlike the conventional methods of forecasting, the proposed framework will not need a significant amount of manual technical analysis and subjective interpretation, but learns the market trends working on the basis of the historical financial information. The objective is to reduce the bias of prediction by human being as well as an effective and stable forecasting model which is computationally effective.

The overall structure of system can be broken down into three large steps, i.e., data preprocessing and feature engineering, deep feature extraction and sequential learning in the shape of hybrid CNN-LSTM model, and a visualization interface that will allow providing an analytical interpretation in real-time.

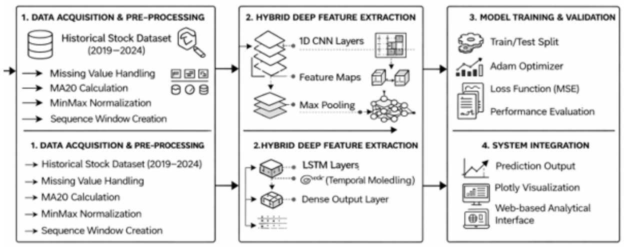

Fig. 1. Architecture of the proposed CNN–LSTM-based stock price prediction model

The initial phase is preprocessing of data. The data on the stocks of a number of technology firms over the years is gathered and refined to get missing or irregular data. Technical features are calculated like the 20-day Moving Average (MA20) to enhance the features space. All the numerical characteristics are scaled with the Min-Max technique so that the gradients can be updated steadily during the training.

The hybrid CNNLSTM structure is the main element of the system. The CNN layers extract short-term trends and local trend characteristics of structured time-series inputs. The extracted feature maps are also not fed directly through dense layers as they are fed to LSTM layers unlike the traditional standalone CNN models. This enables the system to have sequential memory and memorise long-term time dependencies in the stock price movements.

CNN and LSTM are enhanced together, and this ensures that the features are reused successfully, and information is flowing in all directions of the layers. The CNN element is a representation of minor variations as opposed to the LSTM element that is a long term variation. The outcome of this combination is that information loss is minimised and predictive consistency is improved without necessarily incurring a significant increase in the computational complexity. The architecture is neatly designed to be lightweight to the extent that it can be easily deployed and has a good forecasting performance.

The dataset will be split into a training and testing set to cover 2019-2024 stock data in order to guarantee good learning. The Adam optimizer is used to perform model optimization and the Mean Squared Error (MSE) is used to measure performance. The trained model shows a better prediction accuracy when compared to the traditional machine learning methods.

To make it practical and usable, the system contains an interactive visualization module created with the help of Plotly. The dynamic graphs also enable the user to compare the forecasted price and the actual price of the stock hence do intuitive finance analysis. The whole pipeline will be scalable, and it can be integrated into web-based platforms of financial analytics without the need to employ a high-end hardware infrastructure.

The above general conclusions are indicative of the fact that the proposed CNNLSTM model is a balanced model that provides an analytical strength as well as realistic practicality to the real world. The system is also used in making informed decisions that have automated forecasts and accuracy on finance.

5. Methodology

Stocks price forecasting is a multidimensional financial forecasting issue that is subjected to volatility, nonlinearity in relationships and time related market dynamics. The stocks in the technology sector especially show quick changes with economic indicators, world events and investor sentiment. To provide good forecasting, a system must have the ability to learn both the short term fluctuations as well as the long term time dependence. In order to overcome these issues, the hybrid CNN-LSTM based strategy is employed in this paper.

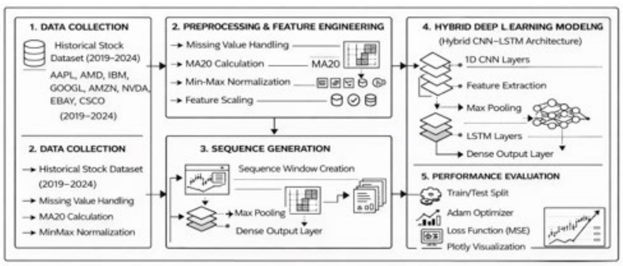

Fig. 2. Proposed research methodology and pipeline of the CNN–LSTM-based stock price prediction system

The suggested methodology will be divided into five key steps, which are data collection, preprocessing and feature engineering, sequence generation, hybrid deep learning modeling, and performance evaluation.

The initial step is the acquisition of data in the past stock history of the leading technological firms such as AAPL, AMD, IBM, GOOGL, AMZN, NVDA, EBAY, and CSCO. The data covers the years 2019-2024 and consists of the attributes of Open, High, Low, Close, and Volume calculated on a daily basis. Financial data usually has blank lines or discrepancies because of non-trading days and data omissions and, therefore, preliminary cleaning is conducted to guarantee continuity and dependability.

The technical features in the preprocessing stage include 20-day Moving Average (MA20), which are determined to improve the space of features. Min-Max normalization is used to do the feature scaling to ensure that the model training is numerically stable. This normalization would lead to a similarity of the input variables, and this would not lead to an instability in gradient in the event of back propagation taking place.

The second one is the transformation of the time-series data into the sequential windows. Rather than entering raw values on a daily basis into the network, sliding window sequences are formed in such a way that every input sample is a constant number of past time steps. The model is able to use this structure to relate very well in terms of time.

The hybrid CNN-LSTM architecture is the main part of the methodology. To begin with, the sequential input data undergoes one-dimensional CNN layers. These layers get localized trends and short run variations of the stock price series. It is the convolution operation that is executed to give resultant feature maps that give rise to the notable trends in the market. Then the max pooling is done to shrink the dimensions and avoid over fitting and leave important features.

The LSTM networks have memory states that enable the system to recall the sequence-based and long-term-dependent patterns unlike the traditional feed-forward neural networks. It is especially applicable in the case of financial forecasting whereby the previous trends determine the future trend of the prices.

Once the features have been learned, this is in turn inputted into a dense layer which will, in turn, give the final estimated price of closing. The data is separated into training and test sets in order to test the performance of generalization. The optimization of the model is done using Adam optimizer and the loss function is Mean Squared Error (MSE) that is used to measure the difference between the predicted price and actual price.

During training, the performance of the model is observed, so that overfitting is avoided, and steady convergence is guaranteed. The trained last model is tested on unknown test data to determine the predictive accuracy. Also, Plotly is used to create interactive visualizations comparing the predicted and actual stock prices to intuitively interpret the performance of the forecasts.

The suggested methodology will combine the power of preprocessing, effective feature extraction, and sequential deep learning to provide a scalable and reliable stock price prediction system to be used in the practical financial analytics use.

6. Result and Discussion

The CNNLSTM model proposed was evaluated to verify whether it could be capable of providing predictions on unknown financial data. Learning rate, batch size, and the count of epochs were also considered in the process of training and the loss function was continuously monitored in the process. The Mean Squared Error (MSE) curve was declining gradually: this was the indicator of stable learning behaviour and effective convergence of model.

The trained network provides the estimated closing prices that are conditioned on learnt time-based patterns with historical information. In contrast to classification, the output layer outputs continuous numerical values, which are future stock prices. To assess the predictive power of the model, it was estimated using unknown data of the years 2019-2024. The results demonstrate the hybrid CNNLSTM model to be effective in the respect of capturing the short-term price fluctuations and the long-term trends.

The curves of loss that were obtained during the training process and validation also exhibited a smooth convergence without a significant difference between the training and validation stages. This implies that overfitting was minimized and the model did not lose its high generalization power.

In order to further examine the model performance, the performance was analyzed by graphical representation of the performance of the models by comparing the predicted and actual stock prices. The prediction curves were very close to the real market trends especially when the prices were moving slowly. It was natural that the deviations were slight in the period when the volatility was extremely high because some unexpected events in the market can happen and sometimes it is impossible to explain them only by referring to the past data.

The suggested model was compared with the classical machine learning approaches of the K-Nearest Neighbors (KNN) and single deep learning architecture. The classical models were also more predictively erroneous particularly on the long-term dependence. On the other hand, CNN-LSTM model recorded low values of MSE and it implies that it is more stable in forecasting.

The efficiency of computation of the model is also another significant observation. Although the architecture incorporates two deep learning elements, it is lightweight and can be trained with the typical computing hardware without any specialized hardware. The generation of predictions takes a few milliseconds, which makes the system suitable to near the real-time financial analysis.

In addition to the accuracy in numbers, automation of stock forecasting minimizes the reliance on subjective interpretation. The analysts can always be different in terms of strategy and experience in the traditional analysis. The specified system may provide the identical outcomes when it is supplied with the similar input information, and, therefore, a reduced degree of fluctuation in the process of the decision-making.

Overall, the experiment outcomes confirm that CNNLSTM model is able to predict the price of stocks in high quality and scale. The effectiveness of the preprocessing stage, hybrid deep learning model utilization, and the structured evaluation demonstrate the possibility of the such systems implementation to the financial analytics environment.

7. Future Work

Even though the suggested CNN-LSTM model delivers good prediction results, additional work can be done. The stock markets are not without various dynamic factors and so the past price data may not necessarily represent the entire market behaviour. In this way, the additional advancement may be performed with the purpose of including some additional sources of information to the accuracy of the predictions.

One of the extensions is the macroeconomic variables that include the interest rates, inflation rates, and the economic indices of other countries. These aspects normally influence the fluctuations of the stock prices significantly. Besides, sentiment analysis, which relies on the articles of news about financial news and information provided by the social media, can also be added to be aware of the psychology of investors. The model might respond to sudden changes in the market more effectively in case they are represented as a combination of numerical time-series data and textual sentiment features.

Another area that is to be enhanced is the integration of attention mechanism into deep learning system. Layers of attention can help the model to identify the time slots, which contribute to the future price dynamics. This can improve the level of forecasting and interpretability. Besides, it is possible to investigate further architectures like Transformer-based models and compare them to the current CNN-LSTM framework.

Another area of future work is the possibility of working on the basis of multi-step forecasting instead of predicting the next time step. The system should be extended to predict prices at longer horizons since it would be more convenient in medium and long-term investment planning. Relationships between two or more stocks can also be studied by portfolio-level prediction in which relations between several stocks are evaluated concurrently.

The visualization methods can be established in a manner that it brings to light the influential characteristics and critical periods that lead to predictions. This would instill confidence in the users and render the financial analysts more transparent in the system.

Finally, the generalization capability will be enhanced due to the expansion of the size of the dataset to include additional companies of other industries and other foreign markets. The model may be trained on continuous basis using new market data to ensure it adapts to fluctuation in the financial conditions.

With these additions, the proposed system can become more realistic and comprehensive financial decision-support system which can be applied practically.

8. Conclusion

In this paper, a CNN-LSTM-type framework hybrid model was created to enhance the accuracy and reliability of stock prices prediction. The system was created to limit reliance on subjective technical analysis by gaining insights into relevant patterns through the historical financial data. The proposed method is capable of capturing short-term fluctuations and long-term temporal dependencies of the market by combining convolutional feature extraction and sequential memory modeling.

The experimental findings indicate that hybrid architecture is more effective than traditional standalone models in prediction stability as well as minimization of errors. The fact that the loss curves converge smoothly and the values of Mean Squared Error are low, says that the model learns more intricate finance behavior without a lot of overfitting. The practical forecasting capability of the system is indicated by the fact that the system can closely track the actual market trends.

However, unlike most research models that stay within the theoretical appraisal, the suggested model offers an end-to-end pipeline that comprises preprocessing, modeling, evaluation, and visualization. The interactive graphical analysis increases usability and the predictions can be interpreted by a simple interpretation by the investors or an analyst. Moreover, the model is computationally efficient, and can be run on common computing equipment, which makes it convenient to be deployed in practice.

The system reduces the variances that occur in the forecast process due to human interpretations by automating the process. This renders the framework reliable when it comes to the subject of financial decision support since the input data will produce the same output at all times. This is particularly critical in circumstances that are volatile in the market where the analysis process can be compromised by emotional inclinations when undertaken manually.

On the whole, the suggested CNNLSTM structure reveals the feasibility of the idea of a deep learning methodology in the field of financial time-series prediction. The experiment proves the fact that the spatial feature extraction and temporal modeling will result in higher predictive power and scalability. The framework can serve as a foundation for developing more advanced intelligent financial analytics systems in the future.

References

[1] Y. Yann LeCun, Y. Yoshua Bengio, and G. Geoffrey Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

[2] S. Sepp Hochreiter and J. Jürgen Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.

[3] A. Alex Graves, A.-r. Mohamed, and G. Geoffrey Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2013, pp. 6645–6649.

[4] Y. Yoon Kim, "Convolutional neural networks for sentence classification," in *Proc. Conf. Empirical Methods Natural Language Processing (EMNLP)*, 2014, pp. 1746–1751.

[5] F. A. Gers, J. Jürgen Schmidhuber, and F. Cummins, "Learning to forget: Continual prediction with LSTM," *Neural Comput.*, vol. 12, no. 10, pp. 2451–2471, 2000.

[6] K. Kim, "Financial time series forecasting using support vector machines," *Neurocomputing*, vol. 55, no. 1–2, pp. 307–319, 2003.

[7] H. Patel, S. Shah, P. Thakkar, and K. Kotecha, "Predicting stock market index using fusion of machine learning techniques," *Expert Syst. Appl.*, vol. 42, no. 4, pp. 2162–2172, 2015.

[8] M. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *Eur. J. Oper. Res.*, vol. 270, no. 2, pp. 654–669, 2018.

[9] Z. Zhang, S. Zohren, and S. Roberts, "Deep learning for financial applications: A survey," *arXiv preprint*, arXiv:2002.05786, 2020.

[10] X. Ding, Y. Zhang, T. Liu, and J. Duan, "Deep learning for event-driven stock prediction," in *Proc. Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2015, pp. 2327–2333.

[11] J. Jason Brownlee, *Deep Learning for Time Series Forecasting*. Melbourne, Australia: Machine Learning Mastery, 2018.

[12] T. Tianqi Chen and C. Carlos Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. ACM SIGKDD Int. Conf. Knowledge Discovery Data Mining*, 2016, pp. 785–794.

[13] D. P. Diederik P. Kingma and J. Jimmy Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2015.

[14] J. James Bergstra and Y. Yoshua Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, pp. 281–305, 2012.

[15] O. Sezer, M. Gudelek, and A. Ozbayoglu, "Financial time series forecasting with deep learning: A systematic literature review," *Appl. Soft Comput.*, vol. 90, p. 106181, 2020.