

Predicting Share Prices Using RNN and Its Variants

Chayan Paul

Associate Professor, Computer Science & Engineering
Department, Swami Vivekananda University, India

Abstract

Calculating share prices is extremely complex because financial markets are intrinsically nonlinear and volatile. The traditional statistical method is not very effective in capturing the complicated patterns that apparently appear in the time series data. Therefore, because they have proven to be highly promising in modeling sequential data by learning temporal dependencies, RNNs along with their advanced variants called LSTM and GRU, have gained popularity for modeling sequential data. This paper will explore the effectiveness of the many uses of RNNs and its derivatives toward price share prediction. The results are indicated in their performance on pre-existing historical stock market data, a comparison for their accurate predictive values, and highlights their attractiveness and limitations. Advanced architectures, such as LSTM, outperform traditional RNN variants by better handling long-term dependencies and overcoming problems with vanishing gradients. The results of this study are analyzed in the context of their implications and future research directions on financial forecasting in the conclusions.

1. Introduction:

The hardest task in financial analytics is that of predicting an accurate share price. Yet, researchers and practitioners continuously return to this task because of the immense economic effect it carries. Various factors affect financial markets, ranging from a set of macroeconomic indicators, company performance metrics, geopolitical events, and

investor sentiment; this makes the problem, amongst others, being associated with forecasting share prices as highly non-linear and overly complex.

Machine learning methods have come out to be quite potent in the capturing of non-linear patterns that exist in data compared to the traditional statistical methods such as with ARIMA models. Of these, especially great promise for time series tasks on forecasting relates to Recurrent Neural Networks whose inherent ability in modeling sequential dependency will enable them to foresee future outcomes rightly.

It analyzes the applicability of RNNs and its two variants, that is, LSTM and GRU in the share price prediction problem. It intends to use historical stock market data so the comparison of performances can evaluate the efficiency of these models. The output sheds light on the effectiveness of deep learning in financial forecasting.

2. Review of Literature

Interest in applying machine learning to financial markets had grown rapidly during the past ten years. Its early methods relied on a much more limited set of techniques, mostly statistical in nature, such as ARIMA, which performed well for linear time series but had a harder time with the inherently non-linear and chaotic nature of financial data.

Introduction of ML introduced several popular practices such as support vector machines, decision trees, and ensemble models, which involve many features in feature engineering and totally ignore the temporal dynamics of stock prices. One of the most significant advancements was neural networks specifically feedforward networks; these can learn non-linear relationships directly from data but they can't model sequential dependencies.

The development of RNNs was a revolutionary point in the history of time series analysis. RNNs are specifically designed for sequential data; it is storing earlier information from the time steps based on a maintained hidden state. However, a traditional RNN suffers from vanishing and exploding

gradient problems that decrease the power of signals for long-term dependencies.

Hochreiter and Schmidhuber suggested LSTM architecture when it comes to dealing with long-term dependencies, including memory cells which are controlled by the use of gating mechanisms in order to control information flow. The variants called GRU are proposed as simpler versions of LSTMs with a similar benefit at reduced complexity. Both such architectures have been used widely and in research studies showing their superiority over traditional approaches.

Despite these developments, there is still a natural hang-up in data preprocessing, model hyperparameter tuning, and choice of evaluation metrics that impact the performance of the models. Furthermore, financial markets are not just governed by economic parameters but depend on several external factors not captured even in historical price data alone, which makes for more work in this area.

3. Methodology:

3.1. Data Collection and Preprocessing

It employed historical stock prices obtained from Yahoo Finance along with other publicly available sources for a chosen subset of companies. Features used are open, high, low, close prices and volume. Interpolation methods are used to fill in the missing values while the data is normalized for numerical stability in training.

3.2. Model Architectures

Three alternate neural network architectures are implemented:

RNN: A basic recurrent architecture with a single hidden layer to lay the baseline.

LSTM: Advanced variant of RNNs which use memory cells and their associated gating mechanisms to handle long-term dependencies.

GRU: A simplified form of LSTM with fewer parameters to make it computationally more efficient.

Implement the model using Python with TensorFlow/Keras on the processed data. Use a rolling-window technique for the data to perform all the experimentations. The experimentations are based on learning rate, batch size, and the number of hidden units.

3.3. Training and Testing

It uses the mean squared error (MSE) loss function and Adam optimizer for training the model. The whole dataset is then split into three data-sets: fitting data with which the models are actually fitted, tuning data (that fine-tune the hyperparameters) and test data for evaluating predictive performance.

The performance of such models can be judged with the help of metrics like the root mean squared error (RMSE) and mean absolute percentage error (MAPE). Qualitative analysis is done in terms of the trends and turning points captured in the prices of the stocks.

3.4. Results

The experimental results show considerable variation between the three models' performance. The traditional RNN performed very poorly on longer-time series traces due to vanishing gradient problems and thus could not achieve optimal levels of predictive accuracy. In contrast, the LSTM model largely outperformed both RNN and GRU, displaying the lowest values for RMSE and MAPE in all datasets. GRU displayed comparable performances compared to the LSTM model but was fairly behind in terms of accuracy, perhaps simply because of its reduced architecture. The qualitative analysis of the predictions indicates that LSTM is much more efficient at capturing trends as well as sudden changes in the stock prices, where RNN oftentimes failed to predict turning points. GRU, though computationally an efficient algorithm, was found unstable at times in its projections, especially during moments of intense market fluctuation. It translates into superiority of the capability of LSTM models in modeling financial time series and, therefore, affirms the status of that preferred architecture for share price prediction tasks.

Conclusion. This research proves the applicability of RNNs

and their derivatives, like LSTM, in predicting share prices. The major findings point towards the fact that complex neural network structures are a must in managing financial data which is inherently non-linear and sequential. Of course, GRUs can be used as a more computational-efficient substitute but lack a little accuracy in comparison with LSTMs. The research also reveals key challenges: its sensitivity to hyperparameters or the necessity to find external data sources in cases of increasing the accuracy of forecasting. Further research may, therefore, be conducted on hybrid models combining LSTMs with some other techniques, such as attention mechanisms, or macroeconomic indicators from an outer source to be used for performance improvement. In conclusion, the results of this study will add to the already growing literature on applications of deep learning in financial forecasting and lay a platform for further research in the same area.

Conclusion

In this chapter, we explored the application of Recurrent Neural Networks (RNNs) and their advanced variants for predicting stock market share prices. Stock market prediction is a highly complex and challenging task because financial markets are influenced by numerous dynamic factors such as economic conditions, investor sentiment, geopolitical events, and historical market trends. Traditional statistical models often struggle to capture the nonlinear and sequential nature of financial time-series data. Deep learning models, particularly RNN-based architectures, offer a powerful alternative by effectively learning patterns from sequential data.

The chapter began with an introduction to time-series forecasting in financial markets, highlighting the importance of historical data such as past stock prices, trading volume, and market indicators. Since stock price data is sequential and time-dependent, models that can learn temporal dependencies are essential. Recurrent Neural Networks were introduced as a suitable approach because they are designed to process sequential inputs and maintain information from previous time steps through internal memory. This capability allows RNNs to capture temporal patterns and dependencies within stock market data.

However, standard RNNs have certain limitations, particularly the vanishing gradient problem, which makes it difficult for them to learn long-term dependencies in long sequences of data. To address this issue, advanced RNN variants such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRU) were discussed. These architectures incorporate gating mechanisms that regulate the flow of information, enabling the model to remember important information for longer periods while discarding irrelevant data. As a result, LSTM and GRU models are widely used in financial forecasting because they provide better performance and stability compared to traditional RNNs.

The chapter also covered the data preprocessing process, which is crucial for building accurate predictive models. Steps such as data cleaning, normalization, feature selection, and sequence generation help transform raw stock market data into a suitable format for training neural networks. Proper preprocessing ensures that the models can learn meaningful patterns without being affected by noise or inconsistencies in the dataset.

Another key aspect discussed was the training and evaluation of the models. The dataset is typically divided into training and testing sets to assess the predictive capability of the models. Performance metrics such as Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE) are commonly used to evaluate how closely the predicted prices match the actual market values. Experimental results generally show that LSTM and GRU models outperform basic RNNs due to their ability to capture long-term dependencies in time-series data.

Furthermore, the chapter emphasized the practical implications of stock price prediction models. Accurate forecasting can assist investors, traders, and financial institutions in making informed decisions regarding buying, selling, and portfolio management. Predictive models can also be integrated into automated trading systems and financial analytics platforms to enhance decision-making processes. However, it is important to recognize that no model can guarantee perfect predictions due to the inherently unpredictable nature of financial markets.

Despite the advantages of RNN-based models, several challenges and limitations remain. Stock markets are influenced by many external factors such as political events, economic policies, global crises, and sudden market sentiment changes that may not be fully captured in historical data. Additionally, deep learning models require large datasets and significant computational resources for training. Overfitting and model generalization also remain important concerns that must be addressed through techniques such as regularization, cross-validation, and hyperparameter tuning.

In conclusion, RNN and its variants such as LSTM and GRU provide powerful tools for modeling and predicting stock price movements, especially when dealing with sequential financial data. These models have demonstrated significant potential in capturing complex temporal patterns that traditional methods often fail to detect. While challenges still exist, the integration of deep learning techniques with financial analysis represents a promising direction for future research. With the continued advancement of machine learning algorithms, increased availability of financial data, and improved computational capabilities, more robust and accurate stock price prediction systems can be developed to support intelligent financial decision-making in modern markets.

References

- Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- K. Cho, B. van Merriënboer, C. Gulcehre, et al., "Learning phrase representations using RNN encoder–decoder for statistical machine translation," in *Proc. EMNLP*, 2014, pp. 1724–1734.
- A. Graves, *Supervised Sequence Labelling with Recurrent Neural Networks*. Berlin, Germany: Springer, 2012.
- I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

T. Fischer and C. Krauss, "Deep learning with long short-term memory networks for financial market predictions," *European Journal of Operational Research*, vol. 270, no. 2, pp. 654–669, 2018.

J. Brownlee, *Deep Learning for Time Series Forecasting. Machine Learning Mastery*, 2018.

A. Selvin, A. Vinayakumar, E. A. Gopalakrishnan, V. K. Menon, and K. P. Soman, "Stock price prediction using LSTM, RNN and CNN-sliding window model," in *Proc. Int. Conf. Advances in Computing, Communications and Informatics (ICACCI)*, 2017, pp. 1643–1647.

M. Nelson, A. Pereira, and R. Oliveira, "Stock market's price movement prediction with LSTM neural networks," in *Proc. Int. Joint Conf. Neural Networks (IJCNN)*, 2017, pp. 1419–1426.

[Z. Zhang, S. Zohren, and S. Roberts, "Deep learning for portfolio optimization," *Journal of Financial Data Science*, vol. 2, no. 4, pp. 8–20, 2020.

R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice*, 2nd ed. Melbourne, Australia: OTexts, 2018.

J. B. Heaton, N. G. Polson, and J. H. Witte, "Deep learning in finance," *Annual Review of Financial Economics*, vol. 11, pp. 1–21, 2019.

S. Siarni-Namini, N. Tavakoli, and A. S. Namin, "The performance of LSTM and BiLSTM in forecasting time series," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, 2019, pp. 3285–3292.

B. McNally, S. Roche, and S. Caton, "Predicting the price of Bitcoin using machine learning," in *Proc. 26th Euromicro Int. Conf. Parallel, Distributed and Network-Based Processing*, 2018, pp. 339–343.

A. Karpathy, "The unreasonable effectiveness of recurrent neural networks," 2015.

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

S. Haykin, *Neural Networks and Learning Machines*, 3rd ed. Upper Saddle River, NJ, USA: Pearson, 2009.

M. Dixon, D. Klabjan, and J. Bang, "Classification-based financial markets prediction using deep neural networks," *Algorithmic Finance*, vol. 6, no. 3–4, pp. 67–77, 2017.

[19] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.

[20] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting*, 2nd ed. New York, NY, USA: Springer, 2016.