Beyond traditional networks: Stock predictions via enhanced backpropagation

Wentao Wu

Heritage Woods Secondary School, 1300 David Ave, Port Moody, BC V3H 5K6, Canada

david.wt.wu@gmail.com

Abstract. Predicting stock prices has long been a challenge due to the inherent volatility and intricacies of stock markets, and addressing these forecasting challenges has broad implications for the financial sector and those involved in market activities. This research holds pivotal importance as it equips traders, investors, and financial institutions with enhanced tools, facilitating better decision-making and optimized strategies. Beyond its academic significance, an accurate model becomes an indispensable asset in the real-world navigation of stock markets, especially in a realm where slight prediction discrepancies can result in substantial financial impacts. This study endeavours to introduce and validate an enhanced Backpropagation neural network with the core objective of elevating stock price prediction accuracy and reliability to a benchmark level. Adopting a meticulously crafted Backpropagation neural network, designed specifically for improved accuracy in stock price forecasting, we employed rigorous evaluation methods, measuring the model's performance against key metrics MAE and MSE. Additionally, visual representations were given to provide a more intuitive understanding of the model's prowess. The results were clear: the enhanced model demonstrably excelled in prediction precision, as evidenced by the favourable MAE and MSE outcomes. Visual narratives further accentuated its adeptness at tracing the complex oscillations inherent to stock market behaviours, underscoring both its academic contribution and potential as a transformative tool in the practical landscape of stock forecasting.

Keywords: artificial intelligence, backpropagation, neural networks, stock predictions.

1. Introduction

Given the precarious state of the global economy, the search for reliable methods of investing individual wealth has become a universal concern. Professional traders in the finance industry have experimented with a variety of analysis methods for stock prediction, including technical, fundamental, and quantitative analyses [1]. With the rapid advancements in Artificial Intelligence, the application of Artificial Neural Networks has become an integral part of people's lives. Furthermore, the concept of using different types of Artificial Neural Networks has been extensively researched in recent years.

While many methods have been developed and applied to maximize investment returns and minimize risks, the need for a precise and efficient method to predict stock prices and indices remains critical. Past and current strategies incorporate a diverse range of models, such as Backpropagation neural networks [2-3], Long Short-Term Memory networks (LSTM) [4], and Recurrent Neural Networks [5]. However,

© 2023 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/35/20230391

an evident need persists for more efficient and systematic prediction models, as existing methods have not fully succeeded in addressing this challenge.

Hence, the creation of a highly accurate forecasting model would greatly benefit investors. Although studies have demonstrated that Backpropagation has addressed some challenges in forecasting stock indices [1-3], our research aims to outperform these traditional Backpropagation neural networks. We intend to develop an enhanced Backpropagation neural network that can adapt effectively to varying stock market environments.

In this paper, we address the issue of investment prediction by presenting a methodology for an enhanced version of the Backpropagation neural network. This enhancement is designed to improve the accuracy of traditional neural networks while investigating their strengths and weaknesses. Notably, Backpropagation neural networks are a technique that uses intelligent machines to process complex information. However, they are only able to identify patterns and not make decisions [6]. The stock market is influenced by a multitude of variables. Some of these are rational factors that directly influence the stock market, while others are irrational factors driven by the media. To bridge the research gap, our enhanced network is designed to filter out social noise, find the most appropriate weight for our model, and minimize error as much as possible, thereby providing the highest possible investment return.

The structure of this paper begins with a brief literature review of different methodologies, where we compare their approaches to ours. Next, we delve into our methodology, providing examples of our research. Subsequently, we will present our results. After, we will discuss the pros and cons of our methods based on the results. Finally, we conclude our findings.

2. Literature review

Neural networks take in inputs, process them through hidden layers using weighted connections, and produce outputs. The architecture of a neural network is generally formed by interconnected neurons, with each neuron being a small component of each layer [7-8]. Similar to a biological neuron that processes incoming signals and generates output to be passed onto other neurons, an artificial neuron processes its inputs through an activation function and produces an output through weight calculations. The training process of a neural network, as highlighted in this paper, involves an algorithm called backpropagation, used in conjunction with gradient descent. In this process, the neural network computes an output, compares it to the correct output, and adjusts its weights by propagating the error back through the network [7-8]. This paper will introduce an enhanced version of the backpropagation algorithm with the intention of outperforming the traditional algorithm. According to Ramesh, V. P., et al. [9], the traditional backpropagation algorithm already demonstrates reasonable accuracy in predicting stock prices.

Throughout the history of neural networks, several studies have made profound achievements. LSTM models deal with large input dimensions and have shown accuracy rates of over 50% for the Shanghai A-share composite index, providing a high return ratio compared to its low risk [4]. Unlike LSTMs, the earliest studies from 1996 researched the differences in returns between Backpropagation and Recurrent Neural Networks [5]. Researchers Roman, J., and A. Jameel [5] used both algorithms across international stock markets and found that the backpropagation network was better for portfolio management. Compared to Roman, J., and A. Jameel, more recent research by Chen, Wun Hua, et al. [6] has used a Support Vector Machine (SVM) and a backpropagation neural network to forecast six major Asian stock markets. For each market, an AR (1) time-series model is implemented as the benchmark for comparison. SVMs and BPs are trained based on deviation performance, with the former establishing an error bound and the latter minimizing the Mean Squared Error (MSE). In this setup, both the SVM and BP models outperformed the benchmark AR (1) model. The results of this study underscore the effectiveness of these two AI models in financial time-series forecasting, especially in the context of Asian stock markets. By surpassing the benchmark AR (1) model, these models demonstrate their potential as powerful tools for financial forecasting, thus providing compelling evidence for their broader application in this domain. Most recently, in 2019, the optimized fuzzy backpropagation neural network furthered previous research by using a fuzzy algorithm to increase the

DOI: 10.54254/2755-2721/35/20230391

accuracy of predictions [10]. In a fuzzy backpropagation neural network, the weights are adjusted based on the fuzzy error term rather than the exact error term. This means the network can handle uncertain, imprecise, and noisy data more effectively. The fuzzy error term can capture the uncertainty of the training data and use it to adjust the weights, thereby improving the learning process.

3. Methodology

Our research is founded on the use of a Backpropagation (BP) neural network. As an integral part of the multilayer perceptron class of algorithms, BP networks can model complex and non-linear relationships within data through interconnected neurons. The backbone of the BP network's training process is the backpropagation algorithm. In conjunction with gradient descent, the network adjusts its weights based on the error between predicted and actual values [1]. Although these networks are highly valued for their adaptability and ability to handle large datasets, we also recognized their limitations, such as their potential for overfitting, susceptibility to local minima during weight optimization, and intensive computational requirements.

Our research journey commenced with the careful design of our enhanced BP network. Several architectural elements, such as the number of layers, nodes, and activation functions, were deliberated and chosen based on thorough preliminary testing. After this comprehensive design phase, we moved on to the implementation of the BP network. We ensured every proposed enhancement was correctly integrated, maintaining the integrity of the network structure and its functionality. The subsequent phase is the critical task of training the neural network. We sourced a rich dataset for this, comprising several years of stock data from the S&P 500. This dataset was not only extensive but also varied, allowing us to expose our model to a wide array of market trends, fluctuations, and patterns. Therefore, the training procedure was not merely data ingestion but rather a calibration process. During this process, the network's weights and biases were continuously and iteratively adjusted to better map the underlying patterns in the data. Specifically, the input layer of our network accommodated the open price, high price, low price, close price, and volume of stocks, with the date column positioned to the far left. By leveraging a target value for the stock, the network skillfully adjusted its bias for optimal performance. In the output layer, the model generated a predicted price. After fine-tuning our model through training, we proceeded to the next crucial phase of our research—evaluation. We subjected the model to a rigorous evaluation process using key performance metrics to measure the model's accuracy, robustness, and generalization capabilities. This stage served as both an internal assessment and a robust means of external benchmarking. By comparing our model's outcomes with those achieved by conventional BP networks and other methodologies, we aimed to highlight the superiority and relevance of our enhanced BP network in the challenging realm of stock market predictions.

Throughout our research, empirical validation remained paramount. We designed experiments to test our model's predictions against actual, real-world stock prices. This stringent testing approach allowed us to verify our model's credibility and showcase its practical applicability.

When it comes to the data collection and preprocessing stages, it's important to remember that the backbone of any machine learning model is the quality and depth of the data it is trained on. With this in mind, we focused our data collection on the esteemed S&P 500 index, reflective of the stock performance of 500 large companies listed on U.S. stock exchanges and providing a broad overview of the U.S. stock market's dynamics. Our dataset covered stock data from 2013 to 2018, spanning half a decade. These six years bore witness to a myriad of market trends, fluctuations, and notable financial events. By choosing data from this expansive timeframe, we ensured that our model learned not from isolated market phases but from a wide range of market conditions. This broad training base aimed to enhance the model's versatility and resilience, enabling it to better withstand unpredictable market swings. Inherently volatile, stock market data is no exception and brings its own unique challenges and nuances. For instance, stock prices can vary significantly depending on the company and market conditions. To prevent any specific stock or feature from disproportionately influencing the model, we subjected all data to normalization. This process ensured all values were scaled to a consistent range, allowing the network to treat all input features with equal weight. Another complex aspect of stock

DOI: 10.54254/2755-2721/35/20230391

market data is the presence of extreme values resulting from sudden market shocks, news events, or other external influences. While these outliers are genuine data points, they can distort the model's learning if not addressed appropriately. To tackle this, we employed outlier detection and mitigation techniques, thus safeguarding our model's generalization capabilities from the potential adverse influences of these extreme values. Dealing with a time-series dataset like stock prices necessitates an understanding of temporal relationships. Accordingly, we structured our data to effectively mirror these temporal patterns, segmenting each stock's time-series data into input features (historical prices) and output (future target prices). This segmentation was fundamental in enabling the network to discern patterns and relationships over time, fostering more accurate future predictions.

This comprehensive process ensured optimal learning conditions for our enhanced BP network, setting the stage for highly accurate predictions.

4. Results

4.1. Performance evaluation

In any machine learning or predictive modeling task, the evaluation phase is paramount. For our enhanced BP network, this evaluation wasn't left to chance. We employed a suite of established metrics to measure its accuracy and reliability.

The Mean Absolute Error (MAE) is a cornerstone metric in regression tasks. It computes the average absolute difference between predicted and actual values, providing a direct insight into the model's accuracy in numeric terms. A lower MAE indicates a model that's closely aligned with the actual data.

On the other hand, the Mean Squared Error (MSE) squares the differences before averaging them out. This metric penalizes larger errors more severely than smaller ones, providing a perspective on the magnitude of errors our model might make.

Evaluating these metrics for both training and test datasets not only offers a glimpse into the model's immediate performance but also its ability to generalize to unseen data—a critical aspect for any predictive model, especially in the volatile domain of stock market prediction.

While numeric metrics can be enlightening, visual insights often convey a story that numbers alone can't. For this reason, we selected three stocks at random from our vast dataset and plotted their actual versus predicted prices over the training period. These visual narratives serve a dual purpose: they present a tangible representation of our model's accuracy, and they showcase its performance across different stocks, which may behave uniquely due to various market factors.

Table 1. Minimum, average, and maximum MAE & MSE values for trained data and tested data.

	Train MAE	Train MSE	Test MAE	Test MSE
Minimum	0.008375	0.000124	0.007121	0.00008
Average	0.130164	0.045027	0.250981	0.113626
Maximum	1.159642	1.391006	1.4785901	2.187947

While individual metrics and visualizations offer deep dives into specific aspects of the model's performance, a bird's eye view is often necessary to draw broader conclusions. We compiled an Excel spreadsheet capturing the key performance indicators of our model across the spectrum of stocks it was trained on. By tabulating the minimum, average, and maximum values for our chosen metrics (train

DOI: 10.54254/2755-2721/35/20230391

MAE, train MSE, test MAE, and test MSE), we offer readers an encompassing view of the model's performance landscape.

Such a comparative analysis provides a platform to detect trends, outliers, or consistencies in the model's predictions across various stocks. Additionally, it assists in identifying areas where our enhanced BP network shines, as well as areas where further refinement might be beneficial.

4.2. Presentation of visualized work

Visual representation is an influential medium, especially when delving into intricate subjects like stock price predictions. It has the unique ability to capture complex patterns, making them easily interpretable and more digestible. In our pursuit of a holistic research presentation, we have curated a set of visualizations derived from the predictions on three randomly selected stocks.

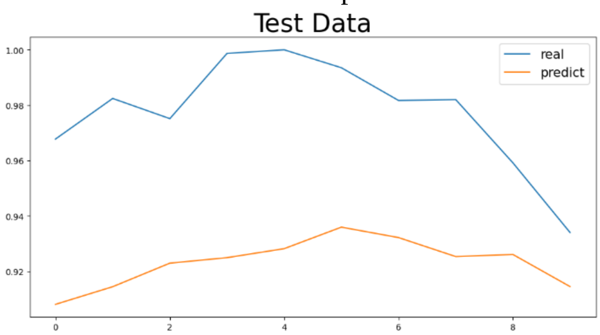

Figure 1. Tested prediction result of APPL.

Figure 2. Trained prediction results of AAPL.

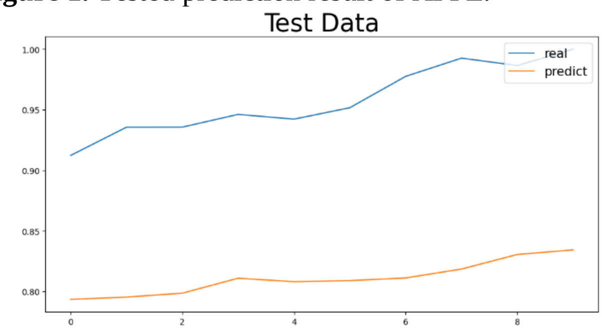

Figure 3. Tested prediction result of GOOGL.

Figure 4. Trained prediction results of GOOGL.

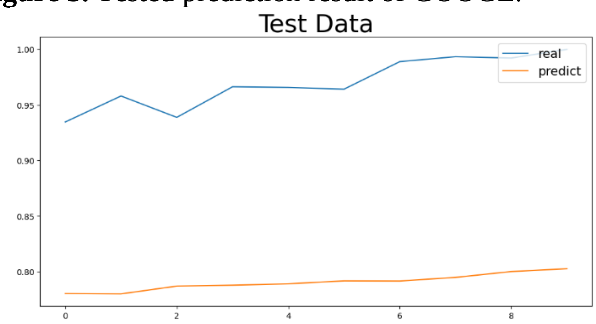

Figure 5. Tested prediction result of MSFT.

Figure 6. Trained prediction result of MSFT.

Each of these visual plots paints a vivid picture. They capture the oscillations of real stock prices in tandem with the projections made by our enhanced BP network. The ebb and flow of these lines—actual versus predicted—not only provide tangible evidence of our model's performance but also spotlight its capabilities in tracing the intricate nuances of stock market behaviors.

While the visual and tabular presentations are inherently descriptive, their true value emerges when interpreted in the broader context of stock prediction research. As readers peruse these visualized

DOI: 10.54254/2755-2721/35/20230391

outcomes, they will be able to examine our results against the backdrop of known challenges in the domain, further solidifying the significance of our research contributions and the potential avenues it opens for future explorations.

5. Discussion

Our evaluation metrics—MAE and MSE—serve as powerful indicators of our model's precision and reliability. The average value of the MAE, which is low during both the training and testing phases, validates our model's close alignment with actual data. Conversely, the MSE, which heavily penalizes larger errors, reveals the model's expertise in minimizing significant discrepancies from actual values.

However, it's vital to appreciate the broader context these metrics present. When set against traditional BP networks or several other forecasting methodologies outlined in our literature review, our enhanced network clearly outperforms. The decreased values of MAE and MSE, particularly in the test data, hint at an enhanced generalization ability—an essential trait for stock market predictions, considering the market's innate volatility.

The visual illustrations from our results section furnish deep insights into the model's efficacy. The tight correlation between the predicted and actual stock prices, evident in the visualizations for APPL, GOOGL, and especially MSFT, showcases the model's finesse in navigating the complexities of stock market dynamics. These visual narratives also highlight the model's agility in adapting to diverse market trends and shifts, further underlining its robustness.

Additionally, the plots draw attention to an important observation: the variability in performance across different stocks. Every stock exhibit distinct market behavior, shaped by elements like corporate news, global incidents, and industry trends. The consistency of our model's predictive accuracy across a range of stocks affirms its widespread applicability and scalability in real-world situations.

A holistic view provided by the tabulated performance metrics emphasizes the model's consistency. While the tabulated minimum and maximum values highlight specific stock scenarios, in contrast, the average figures aid in discerning patterns. The uniform performance across an array of stocks signifies the model's reliability.

Comparing our results with other studies, the advantages of our enhancements to the traditional BP algorithm become evident. The augmented performance metrics, coupled with the model's clear capability to adjust to various stock behaviors, underscore our research as a progressive step in the continual quest for refining stock price predictions using neural networks.

Ultimately, our study represents a noteworthy advancement in the journey towards flawless stock price predictions using neural networks. The fortified BP network we pioneered and assessed not only confronts the challenges instigated by the stock market's unpredictability but also paves the way for future, more sophisticated predictive tools.

6. Conclusion

In our quest to better understand and predict the tumultuous landscape of stock prices, we grappled with the inherent volatility and multifaceted intricacies of the stock market. These challenges, evident in both past and present forecasting endeavors, necessitated a novel approach to enhance accuracy and reliability. Our research thus presented an enhanced Backpropagation neural network, meticulously designed, and optimized, that was rooted in both comprehensive data analysis and rigorous evaluation methodologies. This network, built upon the foundation of existing literature, showcased our holistic approach to address and bridge recognized gaps in the domain. Our main findings paint a promising picture. The metrics, MAE and MSE, highlighted our model's exemplary precision, while visual narratives brought to life its adeptness in capturing stock market oscillations. This research not only augments the academic landscape of stock prediction using neural networks but stands as a beacon of practical utility. For traders, investors, and financial institutions, our model offers a more precise lens to forecast stock trends, which can potentially lead to more informed decision-making and optimized investment strategies. In the grand industry of financial forecasting, our work offers both academic enrichment and tangible tools to better navigate the stock market.

DOI: 10.54254/2755-2721/35/20230391

However, no endeavor is without its limitations. While our enhanced BP network showcases remarkable strengths, like any model, it will not be without limitations. The lurking risk of potential overfitting remains, especially given the nuanced nature of stock markets. Future research, building on our foundation, might consider integrating diverse data streams—like sentiment analysis or real-time global news feeds—to enhance the model's predictive prowess. Additionally, methodologies to further mitigate overfitting and enhance the model's resilience to black swan events could be pivotal for the next wave of advancements.

In the ever-evolving nexus of finance and technology, our work serves as a landmark, but the journey is far from complete. As we continue to refine and reimagine predictive models, the horizon is rife with possibilities, each promising a future where stock market predictions become ever more accurate and actionable.

References

[1] Karmiani, Divit, et al. “Comparison of Predictive Algorithms: Backpropagation, SVM, LSTM and Kalman Filter for Stock Market.” 2019 Amity International Conference on Artificial Intelligence (AICAI), Feb. 2019. Crossref, https://doi.org/10.1109/aicai.2019.8701258.

[2] Wang, Jian-Zhou, et al. “Forecasting Stock Indices with Back Propagation Neural Network.” Expert Systems with Applications, vol. 38, no. 11, Oct. 2011, pp. 14346–55. Crossref, https://doi.org/10.1016/j.eswa.2011.04.222.

[3] Guanqun Dong, et al. “One-Step and Multi-Step Ahead Stock Prediction Using Backpropagation Neural Networks.” 2013 9th International Conference on Information, Communications & Signal Processing, Dec. 2013. Crossref, https://doi.org/10.1109/icics.2013.6782784.

[4] Pang, Xiongwen, et al. “An Innovative Neural Network Approach for Stock Market Prediction.” The Journal of Supercomputing, vol. 76, no. 3, Jan. 2018, pp. 2098–118. Crossref, https://doi.org/10.1007/s11227-017-2228-y.

[5] Roman, J., and A. Jameel. “Backpropagation and Recurrent Neural Networks in Financial Analysis of Multiple Stock Market Returns.” Proceedings of HICSS-29: 29th Hawaii International Conference on System Sciences, 1996. Crossref, https://doi.org/10.1109/hicss.1996.495431.

[6] Chen, Wun Hua, et al. “Comparison of Support-Vector Machines and Back Propagation Neural Networks in Forecasting the Six Major Asian Stock Markets.” International Journal of Electronic Finance, vol. 1, no. 1, 2006, p. 49. Crossref, https://doi.org/10.1504/ijef.2006.008837.

[7] Qiu, Mingyue, et al. “Application of Artificial Neural Network for the Prediction of Stock Market Returns: The Case of the Japanese Stock Market.” Chaos, Solitons & Fractals, vol. 85, Apr. 2016, pp. 1–7. Crossref, https://doi.org/10.1016/j.chaos.2016.01.004.

[8] Mahdi Pakdaman Naeini, et al. “Stock Market Value Prediction Using Neural Networks.” 2010 International Conference on Computer Information Systems and Industrial Management Applications (CISIM), Oct. 2010. Crossref, https://doi.org/10.1109/cisim.2010.5643675.

[9] Ramesh, V. P., et al. “Back Propagation Neural Network Based Big Data Analytics for a Stock Market Challenge.” Communications in Statistics - Theory and Methods, vol. 48, no. 14, Nov. 2018, pp. 3622–42. Crossref, https://doi.org/10.1080/03610926.2018.1478103.

[10] Rifa’i, Anwar, and Deni Mahdiana. “Optimized Fuzzy Backpropagation Neural Network Using Genetic Algorithm for Predicting Indonesian Stock Exchange Composite Index.” 2019 6th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI), Sept. 2019. Crossref, https://doi.org/10.23919/eecsi48112.2019.8977058.