Review: Application of Machine Learning to Investment Portfolios

Yifei Wang

University of Toronto, 27 King's College Circle, Toronto, Ontario M5S 1A1, Canada

yifeiyifei.wang@mail.utoronto.ca

Abstract. This review focuses on the use of three machine learning methods in portfolios using reinforcement learning, recurrent neural networks and random forests. Machine learning algorithms create a model from training or sample data, train themselves by continuously learning from the data, and then adjust their actions based on the insights found to gain the ability to make better predictions and decisions. As today's big data continues to expand and grow, the market demand for machine learning is expected to increase significantly. This is because companies can use it to understand trends in customer behavior and business operation patterns to support the development of new products. Thus, the use of machine learning methods to build an optimal portfolio can help investors minimize risk and maximize returns. Asset management firms enable investors to invest in the best investment opportunities by developing investment plans based on specific client requirements and return expectations.

Keywords: Machine learning, Investment portfolios, Big data.

1. Introduction

The term "portfolio" refers to a group of financial investments, such as stocks, bonds, cash, real estate, works of art, etc. Stocks, bonds, and cash are the most common and core. Diversification is the key to an investment portfolio. Diversification maximizes returns by investing in different areas that react differently to the same investment, while reducing risk by allocating them. It is an investment trust that invests in a predetermined set of stocks and bonds selected by the fund company. The monetary value of each asset affects the risk and rate of return of the portfolio. Building an optimal portfolio can help investors minimize risk and maximize returns. Asset managers enable investors to invest in the best investment opportunities by tailoring investment plans to the specific requirements and return expectations of their clients. The risk preferences and time horizon of clients also impact the portfolio's content dramatically. For instance, an investor who is willing to take risks may invest in an aggressive stock, with high returns and high risk. When constructing a portfolio, investors should take their time horizon into account, similar to their risk tolerance. For example, conservative investors may prefer a portfolio of large-cap stocks, a wide range of bonds, and highly liquid cash equivalents. A specific application of artificial intelligence called machine learning gives machines the capacity to automatically learn from experience and advance without the need for explicit programming. Typically, machine learning algorithms create a model from training data, or sample data. That is to say, machine learning is the ability of computers to train themselves without human intervention by constantly learning from data, and then adjusting their actions accordingly to the insights they discover, thereby gaining the ability to make better predictions and decisions. Machine learning is widely used in various fields, such as finance, medicine and so on. The market demand for machine learning is expected to increase significantly with the continuous expansion and growth of big data nowadays. Because enterprises can understand the trend of customer behavior and business operation mode through it, so as to support the development of new products. Data-driven decisions are increasingly critical in competitive markets. Making decisions that keep businesses ahead of the competition may depend on machine learning, which holds the key to releasing the value of company and customer data. Machine learning is a key component of the business operations of many of today's top corporations, like Facebook, Google, and others.

Portfolio allocation has benefited from the emergence of large-scale portfolio optimization algorithms, as Sarah Perrin demonstrates [1]. The three fundamental models of smart beta portfolios that Sarah looks at are the most diversified portfolio, the risk budget portfolio, and the equal risk allocation portfolio. The development of robo-advisers using machine learning has two key benefits. First, it comprehends investors more thoroughly than conventional asset managers. The second is that it can create an automated rebalancing process and completes the operation in a methodical manner. Tae Kyun Lee's paper proposes financial network indicators for application in global stock market investment strategies [2]. Tae Kyun Le builds directed, and undirected volatility networks of stock markets based on the linkage and correlation of national stock indexes using vector autoregressive models. They studied if network metrics had an effect by using a range of machine learning approaches, such as logistic regression, support vector machines, and random forests as inputs to create strategies. The findings demonstrate the value of the network index as a tool for forecasting the direction and performance of the world stock market. Tae Kyun Lee's study based on financial network indicators to carry out machine learning, to construct the global stock market investment strategies. In Marcos Lopez de Prado's paper, 10 well-known financial applications are used to prove that machine learning has outlived the hype [3]. Marcos showed the case for machine learning to be applied to current financial problems, and machine learning has a good future in finance. Felipe Dias Paiva's research designed a model for intraday trading investment decisions in the stock market [4]. The experimental evaluation of the model is based on the assets of São Paulo Stock Exchange Index for new training and combinatorial optimization of the classification algorithm, with a total of 81 parameter arrangements formulated. Felipe presented two additional models and compared the outcomes with Ibovespa's performance in order to compare the performance of the provided model. For analysis, data from 3716 trading days are used. Along with analysing the impact of brokerage fees on stock purchases and sells on the market, they evaluated the performance of the classifier, the basis of the portfolio, the return and risk of the model, as well as the return and risk of the model. Their methodology still yields significant results despite the demand for transaction value. This study presents a fresh, practical approach to stock price forecasting while also extending the theoretical application of machine learning.

This article mainly reviews the application of machine learning in investment portfolios. It mainly discusses the following three machine Learning methods: Reinforcement Learning, Recurrent Neural Network and Random Forest. It will include an overview of each category and its use in the portfolio, as well as different experimental results and some analysis of this.

2. Using Reinforcement Learning

In reinforcement learning, the intelligent algorithm makes decisions at each stage based on the environment it has seen and the incentives it has received. It learns by making mistakes, with the ultimate objective of selecting the action that would ultimately maximize the total reward. Using a continuous layer, reinforcement learning in finance collects strong information from the original input. Then, to maximize benefits and lower risks, the reinforcement learning model assesses the incoming data using the discovered patterns and methods. In order to address the portfolio management challenge, Zhengyao Jiang suggested a framework for reinforcement learning that is devoid of financial models. The framework consists of four parts: the Ensemble of Identical Independent Evaluators topology, a Portfolio-Vector Memory, An Online Stochastic Batch Learning scheme, and a reward function. The framework is implemented by Convolutional Neural Network, a basic Recurrent Neural Network, and a Long Short-Term Memory. Over the course of a 30-minute trading session, they ran three backtest experiments on the bitcoin market. The three trials of this framework are found to perform better than all other comparative trading algorithms and to get the top three places. Despite the high commission rate, the framework can achieve a return of at least four times within 50 days [5]. According to Zhicheng Wang, current reinforcement learning-based portfolio management models are limited in their ability to balance risk and return. They suggest a deep reinforcement learning technique to improve investing strategies in the paper. This model uses a negative maximum retracement as a return function and macro market circumstances as indicators to lower the risk of market volatility. The combination of classification and correlation can result in a portfolio that follows market trends and strikes a balance between return and risk. The outcomes of research conducted on three well-known stock indices show that the model is much better regarding the risk-return criterion [6]. Bitcoin and other cryptocurrencies are electronic, decentralised replacements for fiat currency. In this paper, Zhengyao Jiang develops a model-free convolutional neural network that receives as input historical prices of a set of financial assets and outputs the set's portfolio weights. The model is made stronger by using 0.7-year pricing data from cryptocurrency exchanges for training, which aims to maximise the cumulative return, which serves as the network's reward function. In the same market, they ran a 30-minute backtest trading experiment and discovered a 10-fold return after 1.8 months. The network, according to Zhengyao Jiang, can be utilised for more than just cryptocurrencies [7].

3. Using Recurrent Neural Network (RNN)

Recurrent neural networks (RNNs) are a type of machine learning in which connections between nodes can establish a loop, allowing the output of certain nodes to influence the following inputs of the same nodes by exploiting timing or time series data. This makes it capable of displaying time-dynamic behaviour. It is suitable for tasks like language translation, speech recognition, and image captioning that don't require segmental connections, such as handwriting recognition or speech recognition. A straightforward regularisation method for recurrent neural networks (RNNs) containing Long Short-Term Memory (LSTM) units was put out by Wojciech Zaremba. Although dropout is the most effective strategy for regularised neural networks, RNNs and LSTMs cannot use it. They presented a simple technique for applying dropout to LSTM and RNNs that can significantly boost performance on a number of issues in various areas. This approach makes it possible for dropout to improve RNNs' application performance [8]. The training of artificial recurrent neural networks (RNNs) using echo state networks and liquid state machines introduces a novel paradigm that trains only one readout from a random-generated RNN. This technique, called reservoir computing, outperforms the original RNNS in many applications and substantially simplifies the actual usage of RNNS. As a result, the original paradigm can be expanded to use various techniques for reservoir training and readout. It can also be used more broadly in various fields, including reservoir adaptation. Mantas Lukoševičius in the paper reviews current methods for generating and tuning reservoirs and training different types of readings, offering a natural conceptual classification of these approaches beyond the realm of conventional reservoir procedures [9]. Parallel Distributed Processing (PDP) architectures show the potential and radical alternative to the traditional paradigm of language processing based on the serial computer model. Learning intricate structural correlations in temporal data poses a danger to PDP systems because, according to automata theory, processing strings from context-free languages (CFL) necessitates the use of a stack or counter storage device. The ability of neural networks to build such a device when learning CFL is not entirely evident, despite the fact that several PDP models have been manually built to imitate one. Recurrent neural networks (RNN) were trained using typical backpropagation training techniques in this study by Paul Rodriguez, who used them to anticipate the next character in a simple deterministic CFL (DCFL). They showed that RNN can be trained to detect the basic DCFL's structure. They also create counter in phase space to ascertain how the network state reflects this structure using dynamic system theory. The findings demonstrate that by demonstrating RNNS in prediction tasks, this approach has the potential to surpass FSM. The classic language processing paradigm based on discrete automata may conform to resource differentiation and computational metaphor, while RNN may not. RNNs could be able to do similarly, however they lack the same mechanical discontinuities [10].

4. Using Random Forest

The third method is Random Forest. Random Forest is also a common machine learning algorithm. During training, several decision trees are built, and the results are obtained by combining the output of various decision trees. It is easy to use and flexible, so it is a learning method that is often used to solve classification and regression problems. The class that many of the trees choose in the classification challenge is the result of the random forest. The average forecast for a single tree is provided for the regression job. Although random forest is often superior to choose tree because it corrects the propensity of overfitting to the training set, gradient enhanced tree still outperforms random forest in terms of accuracy. In applications, random forests are often used as "black box" models because they can produce reasonable predictions on a wide range of data without requiring much configuration. Pushpendu Ghosh uses random Forest and LSTM networks as training methods to analyze their effectiveness in predicting out-of-sample directional movements of S&P 500 stocks in intraday trading from January 1993 to December 2018. In this paper, a multi-feature setting is introduced that includes not only the return closing price relative to, but also the opening price and intraday return. They take Krauss et al.(2017) and Fischer & Krauss(2018) as benchmark. Every day, they short sell the ten stocks with the lowest probability and purchase the ten stocks with the greatest probability. The results show that the daily rate of return provided by the multi-feature setting using LSTM network and random forest is 0.64% and 0.54% respectively. This is better than the benchmark which only includes daily returns relative to the closing price, which is 0.41% for LSTM and 0.39% for Random Forest [11]. In Zheng Tan's study, the robustness of the Random Forest model is evaluated by training on stocks listed on the Chinese stock exchange. Zheng Tan uses both the pure momentum feature field and the fundamental/technical feature space to forecast both short-term and long-term price fluctuations. The data includes every company that was listed on the Chinese stock exchange between February 8, 2013, and August 8, 2017. The Random Forest model is used to forecast the likelihood that each stock will belong to the category with the highest excess return in the upcoming trading session. Each stock is classified into N classes based on its forward excess return. The portfolio components are changed in accordance with the new probabilistic ordering when the stocks chosen to make up the portfolio are held for a period. The findings demonstrate that both feature paradigms have seen substantial excess returns during the last five out-of-sample years. Our findings indicate an inefficient market that is far from equilibrium, even if excess returns have recently reduced in comparison to multi-factor strategies [12]. Random Forest algorithm is a very robust prediction model. It processes huge data sets efficiently and may be applied to classification and regression issues. The accuracy of prediction outcomes is higher with random forest than with decision tree, and the prediction impact is better and simpler to understand.

5. Conclusion

In the first category of machine learning methods, reinforcement learning is used to maximize revenue and reduce risk, and reinforcement learning models are used to evaluate incoming data. Another study on reinforcement learning developed a financial model-free reinforcement learning framework that could achieve at least a fourfold return in 50 days. In addition, a deep reinforcement learning technique is proposed to improve the investment strategy, to improve the risk-return criterion of the model. Also, a model-free convolutional neural network was developed and found to have a 10-fold return after 1.8 months. The second machine learning method, Recurrent Neural Network (RNN), is suitable for tasks such as language translation, speech recognition and image captioning that do not require segmented connections. A direct regularization method for recurrent neural networks (RNNs) with long short-term memory units is proposed to improve the performance of RNN applications. In addition, methods for adjusting reservoirs and training different types of readings have the effect of extending reservoir calculations. Another study showed that the method has the potential to outperform FSM by demonstrating RNNS in a prediction task. The last machine learning method is Random Forest, which can produce reasonable predictions on a wide range of data without much configuration. Some studies have used random forest and LSTM networks as training methods and introduced a multi-feature setting, followed by stock buying and selling, which has shown better results than a benchmark that only includes daily returns (relative to closing prices). The robustness of Random Forest is also evaluated and considered to be a very robust prediction model. In conclusion, these studies show that machine learning can optimize portfolios, and predictions from machine learning can help overcome many problems. In future research, one can explore the role of a wider range of optimizers and make improvements that may result in better measures that increase portfolio returns and reduce risk.

References

[1] Tae Kyun Lee, Joon Hyung Cho, Deuk Sin Kwon, So Young Sohn, Global stock market investment strategies based on financial network indicators using machine learning techniques, Expert Systems with Applications, Volume 117, 2019, Pages 228-242, ISSN 0957-4174, https://doi.org/10.1016/j.eswa.2018.09.005.

[2] López de Prado, Marcos and López de Prado, Marcos, Ten Applications of Financial Machine Learning (September 22, 2019). Available at SSRN: https://ssrn.com/abstract=3365271.

[3] Felipe Dias Paiva, Rodrigo Tomás Nogueira Cardoso, Gustavo Peixoto Hanaoka, Wendel Moreira Duarte, Decision-making for financial trading: A fusion approach of machine learning and portfolio selection, Expert Systems with Applications, Volume 115, 2019, Pages 635-655, ISSN 0957-4174, https://doi.org/10.1016/j.eswa.2018.08.003.

[4] Perrin, S. and Roncalli, T. (2020). Machine Learning Optimization Algorithms & Portfolio Allocation. In Machine Learning for Asset Management, E. Jurczenko (Ed.). https://doi.org/10.1002/9781119751182.ch8.

[5] Jiang, Zhengyao, Dixing Xu, and Jinjun Liang. "A deep reinforcement learning framework for the financial portfolio management problem." arXiv preprint arXiv:1706.10059 (2017).

[6] Wang, Z., Huang, B., Tu, S., Zhang, K., & Xu, L. (2021). DeepTrader: A Deep Reinforcement Learning Approach for Risk-Return Balanced Portfolio Management with Market Conditions Embedding. Proceedings of the AAAI Conference on Artificial Intelligence, 35(1), 643-650. Retrieved from https://ojs.aaai.org/index.php/AAAI/article/view/16144.

[7] Z. Jiang and J. Liang, "Cryptocurrency portfolio management with deep reinforcement learning," 2017 Intelligent Systems Conference (IntelliSys), 2017, pp. 905-913, doi: 10.1109/IntelliSys.2017.8324237.

[8] Zaremba, W., Sutskever, I., & Vinyals, O. (2014). Recurrent neural network regularization. arXiv preprint arXiv:1409.2329.

[9] Mantas Lukoševičius, Herbert Jaeger, Reservoir computing approaches to recurrent neural network training, Computer Science Review, Volume 3, Issue 3, 2009, Pages 127-149, ISSN 1574-0137, https://doi.org/10.1016/j.cosrev.2009.03.005.

[10] Rodriguez, Paul, Janet Wiles, and Jeffrey L. Elman. "A recurrent neural network that learns to count." Connection Science 11.1 (1999): 5-40.

[11] Pushpendu Ghosh, Ariel Neufeld, Jajati Keshari Sahoo, Forecasting directional movements of stock prices for intraday trading using LSTM and random forests, Finance Research Letters, Volume 46, Part A, 2022, 102280, ISSN 1544-6123, https://doi.org/10.1016/j.frl.2021.102280.

[12] Zheng Tan, Ziqin Yan, Guangwei Zhu, Stock selection with random forest: An exploitation of excess return in the Chinese stock market, Heliyon, Volume 5, Issue 8, 2019, e02310, ISSN 2405-8440, https://doi.org/10.1016/j.heliyon.2019.e02310.