A Survey of Deep Reinforcement Learning in Financial Markets

Ying Yu

Antai College of Economics & Management, Shanghai Jiao Tong University, 200240, China

yuying_mel@sina.com

Abstract. This paper surveys the application of reinforcement learning (RL) in stock price prediction, highlighting its potential and limitations. We explore how RL can be used to optimize trading strategies, manage investment risks, find arbitrage opportunities, and predict trends. The review classifies research objects and methods based on data frequency (high/non-high) and target (forecast/trading strategy). We analyze various asset classes (stocks, forex, etc.) and models (RL, neural networks, LSTMs) employed in previous works. Key findings suggest that RL offers advantages over traditional models by adapting to complex market dynamics, and that incorporating sentiment analysis can further enhance its effectiveness. We identify promising avenues for future research, including hybrid models, deeper sentiment integration, and improved risk management. Overall, the paper concludes that RL holds significant promise for transforming financial forecasting, leading to more accurate and adaptable decision-making tools.

Keywords: Reinforcement learning, stock price prediction, financial forecasting, sentiment analysis, deep learning, machine learning, artificial intelligence

1 Introduction

Reinforcement learning is a machine learning method that learns by having an agent take actions in its environment and adjust its policy based on feedback. In stock price prediction, reinforcement learning can be used to build an intelligent agent that can take actions in a stock market environment (such as buying or selling a stock) and adjust its strategy based on market feedback (such as changes in stock prices).

First, reinforcement learning can help us optimize trading strategies to maximize returns. By training an intelligent agent, we can make it learn to take optimal actions under different market conditions. This may include determining when to buy, sell or hold stocks and how to allocate weights in the portfolio.

Next, reinforcement learning can help us better manage investment risks. By training an intelligent agent, we can make it learn to act conservatively or aggressively under different market conditions to reduce potential losses. This may include setting stop-

K. Subramaniyam et al. (eds.), *Proceedings of the 3rd International Academic Conference on Blockchain, Information Technology and Smart Finance (ICBIS 2024)*, Atlantis Highlights in Computer Sciences 21,
https://doi.org/10.2991/978-94-6463-419-8_24

loss and take-profit points, as well as adjusting portfolio weightings based on market fluctuations.

Third, reinforcement learning can help us find arbitrage opportunities between futures contracts with different expiration dates. By training an intelligent agent, we can make it learn to perform intertemporal arbitrage under different market conditions to achieve stable returns.

Finally, reinforcement learning can help us predict stock market trends. By training an intelligent agent, we can make it learn to analyze historical data and predict future price movements. This may include using technical indicators, fundamental data, and other relevant information to predict market trends.

In summary, the application of reinforcement learning in stock price prediction has broad potential. However, due to the complexity and uncertainty of the stock market, reinforcement learning models may require large amounts of data and computing resources for training, and may need to be continuously adjusted and optimized to adapt to the changing market environment.

2 Classification of Research Objects

Table 1. Different Research Objects

Data	Objects	
	Forecast	Trading Strategy
Non high-frequency data	[21-24]	[5-8][9-18]
High Frequency Data	[1-4]	[19-20]

2.1 Criteria

In this section, two independent and different criteria would be used to divide research objects into different types: 1) Input Data: daily/monthly/hourly bars and high frequency limit order book data. The bars data are used for the purpose of non high frequency trading. 2) Objects: Forecast and trading strategies. For trading strategy, the aim is to discover investment decisions. For forecast, the results are based on forecasts, but also embed other information, such as the certainty of the forecast (how much should we invest) that influence decisions.

2.2 The Classification

Based on the appeal classification standard, we give the classification in Table 1. The meaning of each class is as follows:

2.2.1Non-high Frequency Data & Forecast.

This type is using non-high frequency data to predict market return or direction. References [21-24] belong to this type. Based on daily historical data, reference [24] developed algorithm to the classification and prediction of stock price patterns. Besides the daily price data, other information also be used as input for the forecast. Reference [21] use daily sentiment data to predict stocks price moves. Reference [22] use internet users attention level to predict the stock price. Reference [23] extract data and trading data from news to predict stock movements.

2.2.2Non-high Frequency Data & Trading Strategy.

This type is using non-high frequency data to discover trading strategies. References [5-8][9-18] belong to this type. With daily data as input, reference [6][9-13] developed models for single stock trading. Reference [14] develop strategies for futures trading. Reference [16] built strategies for trading multiple stocks. Reference [17] built strategies for trading multiple cryptocurrencies. Reference [18] developed finance portfolio optimizing strategies for stock trading.

2.2.3High Frequency Data & Forecast.

This type is using high frequency data to predict market return or direction. References [1-4] belong to this type. Reference [1] use the information in Limit Order Book to forecast stock prices' short term movement. Reference [2] forecast stock price based on order flow information. Reference [3] forecast stock price move direction based on Limit Order Book information. Reference [4] forecast the stock price's quantile based on limit order book information.

2.2.4High Frequency Data & Trading Strategy.

This type is using high frequency data to discover trading strategies. References [19-20] belong to this type. Reference [19] developed high frequency market making strategy. Reference [20] did research for large tick asset's market making strategies.

3 Classification of Research Methods

The classificaiton of research methods is shown in Table 2.

Table 2. Different Research Methods

Model	Asset Class				
	Stocks	FX	Fixed Income	Commodities	Cryptocurrencies
RL	[5-13][15-16][18-21]	[5][14]	[5][7]	[12][14]	
NN	[3-4][11][22][24]				[17]
LSTM	[1-2][4][10-11][22-23]				

3.1 Criteria

In this section, two independent and different criteria would be used to divide rsearch objects' asset classes into different types: 1) Asset classes, which includes: A) Stocks. This class includes single stock, multiple stocks, stock index, and stock index futures.

b) Foreign exchange rates. It's the currencies rate and its futures contracts. c) Fixed Income. T-Bills is in this class. d) Commodities. It's commodities futures contracts. e) Cryptocurrencies. The most valued cryptocurrencies are included in this class. 2) Models, which includes: A) Reinforcement Learning, a type of machine learning inspired by how animals learn through trial and error. B) Neural Network, a powerful machine learning model inspired by the structure and function of the human brain. C) LSTM, i.e., Long Short-Term Memory, a type of recurrent neural network (RNN) specifically designed to address the vanishing gradient problem.

3.2 Explanation of Different Types

Ref	Methods	Assets Classes
5	RRL	U.S. Dollar/British Pound Foreign Exchange; S&P 500 and T-Bill
6	RRL	S&P 500 stock index
7	RRL	S&P 500 / TBill; macroeconomic data
8	QL	German stock index DAX
9	QL	Single stock
10	DQL, LSTM	S&P500 ETF
11	RNN, LSTM	7, 000 US-based stocks
12	DL, RL	China CSI300 stock index future contracts; Silver and Sugar futures contracts
13	DNN, DQN	S&P 500 index, HSI index, STOXX 50 index and KOSPI index
14	DRL	50 very liquid futures contracts of across different asset classes,
15	DQN	30 stocks presenting
		diverse characteristics (sectors, regions, volatility, liquidity, etc.
16	DRL	Dow Jones 30 constituent stocks
17	CNN	12 most-volumed cryptocurrencies
18	RL, DL	stocks
19	QL	stocks
20	RL	three representative large tick stocks MSFT, INTC and GE
21	LR, SVM, DL	four top stock companies; intensive dataset of tweets
22	RNN, LSTM and GRU	the data of users and their self-selected stocks; SSE 50 constituent stocks
23	LSTM	stocks
24	BP NN	3 stocks in chinese market

4 Final Remarks and Conclusions

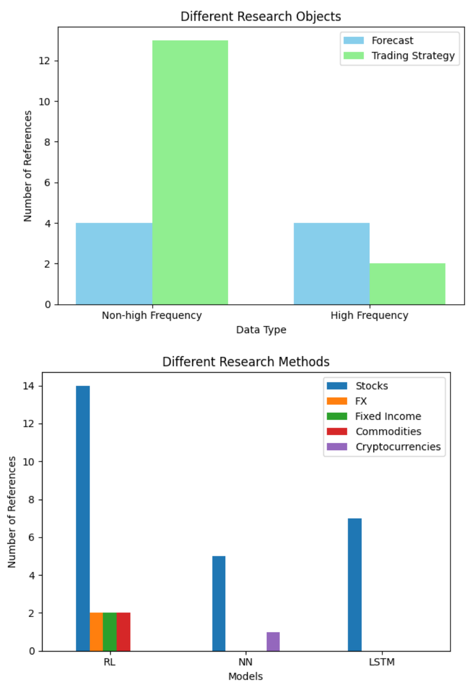

Fig. 1. Distribution of literatures on research objects and methods

Figure 1 provides the distribution of literatures on research objects and methods. We can see that most studied object is on non-high frequency data with trading strategy, and most used method is RL on stocks. This review paper concludes that reinforcement learning (RL) offers substantial improvements for financial market forecasting, outperforming traditional models by adapting to complex, dynamic data patterns. Including sentiment analysis further enhances RL models' predictive capabilities. The paper serves as a research guide, suggesting future exploration in hybrid models, sentiment integration, and risk management.

References

1. Tsantekidis, A., Passalis, N., Tefas, A., Kanniainen, J., Gabbouj, M., & Iosifidis, A. (2017, July). Forecasting stock prices from the limit order book using convolutional neural networks. In 2017 IEEE 19th conference on business informatics (CBI) (Vol. 1, pp. 7-12). IEEE.
2. Sirignano, J., & Cont, R. (2021). Universal features of price formation in financial markets: perspectives from deep learning. In Machine Learning and AI in Finance (pp. 5-15). Routledge.

3. Zhang, Z., Zohren, S., & Roberts, S. (2018). Bdlob: Bayesian deep convolutional neural networks for limit order books. arXiv preprint arXiv:1811.10041.
4. Zhang, Z., Zohren, S., & Roberts, S. (2019). Extending deep learning models for limit order books to quantile regression. arXiv preprint arXiv:1906.04404.
5. Moody, J., & Saffell, M. (2001). Learning to trade via direct reinforcement. IEEE transactions on neural Networks, 12(4), 875-889.
6. Moody, J., & Saffell, M. (1998). Reinforcement learning for trading. Advances in Neural Information Processing Systems, 11.
7. Moody, J., Wu, L., Liao, Y., & Saffell, M. (1998). Performance functions and reinforcement learning for trading systems and portfolios. Journal of forecasting, 17(5-6), 441-470.
8. Neuneier, R. (1995). Optimal asset allocation using adaptive dynamic programming. Advances in neural information processing systems,
9. Bertoluzzo, F., & Corazza, M. (2012). Testing different reinforcement learning configurations for financial trading: Introduction and applications. Procedia Economics and Finance, 3, 68-77.
10. Chen, L., & Gao, Q. (2019, October). Application of deep reinforcement learning on automated stock trading. In 2019 IEEE 10th International Conference on Software Engineering and Service Science (ICSESS) (pp. 29-33). IEEE.
11. Dang, Q. V. (2019, December). Reinforcement learning in stock trading. In International conference on computer science, applied mathematics and applications (pp. 311-322). Cham: Springer International Publishing.
12. Deng, Y., Bao, F., Kong, Y., Ren, Z., & Dai, Q. (2016). Deep direct reinforcement learning for financial signal representation and trading. IEEE transactions on neural networks and learning systems, 28(3), 653-664.
13. Jeong, G., & Kim, H. Y. (2019). Improving financial trading decisions using deep Q-learning: Predicting the number of shares, action strategies, and transfer learning. Expert Systems with Applications, 117, 125-138.
14. Zhang, Z., Zohren, S., & Stephen, R. (2020). Deep reinforcement learning for trading. The Journal of Financial Data Science.
15. Théate, T., & Ernst, D. (2021). An application of deep reinforcement learning to algorithmic trading. Expert Systems with Applications, 173, 114632.
16. Yang, H., Liu, X. Y., Zhong, S., & Walid, A. (2020, October). Deep reinforcement learning for automated stock trading: An ensemble strategy. In Proceedings of the first ACM international conference on AI in finance (pp. 1-8).
17. Jiang, Z., & Liang, J. (2017, September). Cryptocurrency portfolio management with deep reinforcement learning. In 2017 Intelligent systems conference (IntelliSys) (pp. 905-913). IEEE.
18. Zhang, Z., Zohren, S., & Roberts, S. (2020). Deep learning for portfolio optimization. The Journal of Financial Data Science.
19. Lim, Y. S., & Gorse, D. (2018, April). Reinforcement learning for high-frequency market making. In ESANN 2018-Proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (pp. 521-526). Esann.
20. Wang, Y. (2019). Electronic market making on large tick assets. The Chinese University of Hong Kong (Hong Kong).
21. Maqsood, H., Mehmood, I., Maqsood, M., Yasir, M., Afzal, S., Aadil, F., ... & Muhammad, K. (2020). A local and global event sentiment based efficient stock exchange forecasting using deep learning. International Journal of Information Management, 50, 432-451.

22. Liu, K., Zhou, J., & Dong, D. (2021). Improving stock price prediction using the long short-term memory model combined with online social networks. Journal of Behavioral and Experimental Finance, 30, 100507.
23. Akhtar, M. M., Zamani, A. S., Khan, S., Shatat, A. S. A., Dilshad, S., & Samdani, F. (2022). Stock market prediction based on statistical data using machine learning algorithms. Journal of King Saud University-Science, 34(4), 101940.
24. Zhang, D., & Lou, S. (2021). The application research of neural network and BP algorithm in stock price pattern classification and prediction. Future Generation Computer Systems, 115, 872-879.