Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 14, No. 1, 2024

Deep Reinforcement Learning-based Algorithmic Optimisation and Risk Management for High Frequency Trading

Fengyuan Liu [1, *], Yuan Tian[2]

[1] Shandong University of Finance and Economics, Jinan, Shandong, 250220, China
[2]School of Journalism and Communication of Southwest University, Chongqing, 311401, China

[*] **Corresponding author**: Fengyuan Liu

Abstract: This paper reviews the current status and challenges of Deep Reinforcement Learning (DRL)-based algorithm optimisation and risk management for high-frequency trading. By analysing the potential application of Deep Reinforcement Learning in high-frequency trading, its unique advantages in algorithm optimisation, trading decision-making and risk management are discussed. Although DRL demonstrates the ability to make self-adaptive and dynamic decisions in complex market environments, it still faces many challenges such as insufficient real-time algorithmic performance, data sparsity, model overfitting, and risk management complexity in practical applications. This paper summarises the main findings of the current research and proposes directions for future research, suggesting that the application of DRL in high-frequency trading can be further enhanced by improving the algorithmic structure, dealing with data sparsity, and optimising risk management strategies.

Keywords: Deep Reinforcement Learning; High Frequency Trading; Algorithm Optimisation; Risk Management; Multi-Intelligent Body System.

1. Introduction

In recent years, with the rapid development of the financial markets and the continuous progress of technology, high-frequency trading (HFT) has become an important part of the financial markets. High-frequency trading executes a large number of trades in a very short period of time through high-speed computer programmes in order to obtain small price differences and thus achieve high returns. However, as market competition intensifies and the trading environment becomes more complex, HFT faces increasing challenges [1]. Traditional high-frequency trading algorithms have gradually revealed some limitations in responding to market volatility, handling large amounts of data and managing trading risks. Therefore, how to optimise high-frequency trading algorithms to improve trading efficiency and reduce risk has become a widely concerned research topic in both academia and the industry.

Deep Reinforcement Learning (DRL), as an important branch of Artificial Intelligence, has gradually emerged in the financial field in recent years, especially in the application of high-frequency trading. Deep Reinforcement Learning combines the advantages of Deep Learning and Reinforcement Learning. By simulating the trading environment, the intelligent body can independently learn the optimal trading strategy and dynamically adjust it in the complex and changing market environment [2]. Compared with traditional rule-based trading strategies, Deep Reinforcement Learning is better able to adapt to non-linear and high-dimensional data, with stronger generalisation and adaptability. This technique not only helps optimise the performance of high-frequency trading algorithms, but also provides new ideas for risk management.

Although deep reinforcement learning shows great potential in high-frequency trading, current research is still in the initial exploration stage. Existing research focuses on the optimisation of algorithms and the formulation of trading strategies, yet deep reinforcement learning algorithms face a series of challenges in practical applications [3]. For example, the sparsity and noise problems of market data, the stability and robustness of the algorithms, and the effectiveness of risk management all require further in-depth research. Therefore, this review aims to systematically sort out the current status of deep reinforcement learning application in high-frequency trading, analyse its advantages and shortcomings in algorithm optimisation and risk management, and explore possible future development directions.

2. Reinforcement Learning Overview

Deep Reinforcement Learning (DRL) is an artificial intelligence approach that has emerged in recent years, combining the advantages of reinforcement learning and deep learning, and showing strong potential for decision-making problems in complex environments. First of all, the basic concept of reinforcement learning, as an important machine learning paradigm, is that by interacting with the environment, an intelligent body (agent) learns policies through trial and error in order to maximise the accumulated long-term rewards [4]. In this process, the agent receives state information from the environment and selects an action based on the current policy, which acts on the environment to provide a new state and immediate reward. The core goal of reinforcement learning is to find an optimal policy such that the action taken in any state maximises the expected future reward.

Deep learning, on the other hand, performs automated feature extraction and pattern recognition on data through multi-layer neural networks, and has significant advantages especially when dealing with high-dimensional data. Deep Reinforcement Learning thus emerged, which enables intelligences to process complex, high-dimensional state spaces by introducing deep neural networks into reinforcement learning frameworks. This is exemplified by the Deep Q-Network (DQN) in Deep Reinforcement Learning, which uses neural networks to approximate the Q-value function in Reinforcement Learning, thus enabling the processing of complex decision-making problems without the need to explicitly model the environment. In addition, Policy Gradient methods and Actor-Critic models are also common algorithms used in deep reinforcement learning to improve the performance of intelligences by continuously updating their policies and value estimates [5].

In the financial field, the application of deep reinforcement learning is gradually gaining attention, especially in the optimisation of high-frequency trading algorithms and risk management. High-frequency trading requires real-time processing of large amounts of market data to make fast and accurate trading decisions. Traditional rule-based strategies often lack flexibility and adaptability in dealing with the rapidly changing market. Deep reinforcement learning enables trading strategies to be dynamically adjusted with a certain degree of adaptivity through simulation training of historical data and market environment. Intelligentsia are able to learn potential patterns in the market during the training process and are still able to make effective trading decisions in the face of unknown market conditions [6].

In addition, the application of deep reinforcement learning in the financial market is not only limited to high-frequency trading, but also includes asset allocation, risk management, hedging strategies and many other aspects. By introducing deep reinforcement learning models, financial systems can better handle uncertainty and non-linear problems, thus improving the profitability and stability of trading. For example, in risk management, deep reinforcement learning can help construct dynamic risk control strategies, enabling the system to flexibly adjust risk exposure under different market conditions, thereby effectively reducing potential losses from extreme market volatility.

Overall, deep reinforcement learning, as a technology that combines the advantages of deep learning and reinforcement learning, shows great potential for application in the financial field, especially in high-frequency trading. Although there are still some challenges, deep reinforcement learning is expected to play a more important role in the financial market in the future as the technology continues to progress and application scenarios are gradually expanded.

3. Status and challenges of high-frequency trading

High-Frequency Trading (HFT) is a type of trading that relies on high-speed computer technology to make quick profits based on small price fluctuations by placing a large number of orders in a very short period of time. Since the late 1990s, high-frequency trading has gradually developed into an important part of the financial market. With the continuous progress of computer technology and communication technology, high-frequency trading has significantly increased its influence and trading volume in the market. The core of high-frequency trading lies in capturing market opportunities and conducting buying and selling operations in milliseconds or even microseconds through highly efficient algorithms and advanced hardware equipment, thus achieving high-frequency short-term investment returns.

High-frequency trading relies on a variety of complex strategies and techniques, the most common of which include Market Making, Arbitrage and Trend Following. Market making strategies profit from the bid-ask spread by simultaneously quoting bid and ask prices; arbitrage strategies take advantage of price differences between markets; and trend following strategies attempt to capture and profit from short-term trends in price fluctuations. In order to implement these strategies, high-frequency trading relies on ultra-low latency trading systems, complex mathematical models and powerful data processing capabilities [7]. At the heart of these technologies is the ability to react quickly to market changes and execute orders in a timely manner.

Despite its remarkable success in the financial markets, high-frequency trading faces a number of challenges and risks. Firstly, high-frequency trading needs to process huge amounts of market data and make trading decisions in a very short period of time, which puts high demands on the efficiency and accuracy of the algorithms. Second, the complexity and unpredictability of the market environment makes high-frequency trading vulnerable to unexpected events or extreme market conditions, which can lead to significant trading risks. In addition, due to the intense competition in high-frequency trading, the problem of homogenisation of trading strategies has become increasingly serious, which makes it difficult for market participants to gain a sustainable competitive advantage. Finally, the constant tightening of regulation and market changes are also placing new demands on high-frequency trading, and traders must continually update their strategies to adapt to new market rules and environments [8].

The limitations of current high-frequency trading algorithms are also becoming apparent. Despite their success in the past, traditional high-frequency trading algorithms have begun to show a lack of adaptability as market competition and data complexity have increased. These algorithms typically rely on fixed rules and models, and struggle to respond effectively to the non-linear and non-stationary nature of the market. In addition, models that rely excessively on historical data often show poor robustness in the face of sudden market changes, leading to greater risk in the practical application of trading strategies [9]. At the same time, high-frequency trading algorithms also face challenges when dealing with risk management, and how to effectively control risks while pursuing high returns has become an important issue in the development of high-frequency trading.

Overall, high-frequency trading plays an important role in the financial market, but the challenges and limitations it faces should not be ignored. With the continuous development of the market and technological advances, how to optimise high-frequency trading algorithms and effectively manage trading risks will be a key issue to be addressed in future research and practice.

4. Deep Reinforcement Learning in High Frequency Trading

In specific use cases, deep reinforcement learning has shown its great potential in high-frequency trading. For example, some studies have used models such as Deep Q Network (DQN) and Actor-Critic to develop systems that can automatically adjust buy and sell strategies in high-frequency trading. These systems are able to process and analyse market data in a very short period of time and make trading decisions quickly, thus achieving higher rates of return amidst market volatility. In addition, there are studies that improve overall trading performance and risk control by introducing multi-intelligence systems that allow multiple intelligences to work together under different market conditions.

Deep reinforcement learning is particularly compelling in trading decisions. Deep Reinforcement Learning is able to adapt more flexibly to the non-linear and non-stationary nature of markets than traditional models based on statistics and economic theory. This allows intelligences to make relatively robust decisions in the face of transient market fluctuations and extremes. For example, intelligences can automatically adjust the frequency and size of trades when market prices fluctuate dramatically, thereby reducing potential losses. At the same time, the adaptive nature of deep reinforcement learning also enables it to choose different trading strategies under different market conditions to optimise the overall investment portfolio [10].

In addition, the application of deep reinforcement learning in market microstructure modelling shows great potential. Market microstructure research is concerned with the process of order generation, execution and trade price formation in the market. Through deep reinforcement learning, the behavioural patterns of different market participants can be simulated and analysed to predict short-term changes in market prices. This is particularly important for high-frequency trading, as capturing subtle changes in the microstructure can lead to significant trading advantages. Some studies have optimised trading decisions by constructing market simulation environments in which intelligences are trained to learn more granular market patterns.

Overall, the application of deep reinforcement learning in high-frequency trading has demonstrated significant benefits, but it also faces challenges. How to further improve the robustness and efficiency of the algorithm and effectively manage risks in real trading remains an important direction for future research. With the continuous progress of technology, deep reinforcement learning is expected to play an increasingly important role in the field of high-frequency trading.

5. Risk management in high-frequency trading

High Frequency Trading (HFT), as one of the most agile and fastest trading methods in the financial markets, is accompanied by a variety of risks, despite being able to achieve significant gains in a short period of time. The types of risks associated with HFT are varied and complex, including market risk, liquidity risk, operational risk and technical risk. Market risk arises from extreme price fluctuations, especially when the market moves dramatically, and failure to adjust positions in a timely manner can lead to significant losses. Liquidity risk, on the other hand, arises from the inability to execute large orders at expected prices due to insufficient market depth, which may trigger slippage or trade failures. Operational risk involves incorrect execution of trading strategies or technical failures, while technical risk includes network delays, system crashes or algorithmic errors [11]. The high incidence and suddenness of these risks make high-frequency trading a huge challenge while pursuing high returns.

Currently, risk management methods in high-frequency trading rely on pre-set risk control mechanisms and real-time monitoring systems. These methods usually include setting stop-loss points, position limits, and risk assessment using risk indicators (e.g. VaR, Value at Risk). Through these measures, traders can limit losses on individual trades to a certain extent and maintain a manageable risk exposure for the portfolio as a whole. In addition, real-time monitoring systems can quickly identify and respond to unusual trading activity or sudden changes in market conditions. However, these traditional risk management approaches are often rigid and lagging in the face of highly dynamic and uncertain market environments, making it difficult to quickly adjust to unexpected market risks.

The potential of Deep Reinforcement Learning (DRL) in risk management is gradually emerging. Compared with traditional methods, deep reinforcement learning is able to adjust trading strategies in real time to cope with changing market environments through continuous learning and optimisation. Deep Reinforcement Learning intelligences can continuously receive feedback from the market during the training process and update their strategies on this basis, thus making more flexible and precise risk control decisions in the face of market volatility. For example, the intelligent body can dynamically adjust position size, trading frequency and stop-loss points based on real-time market data to minimise potential losses. In addition, Deep Reinforcement Learning is able to consider multiple risk factors, such as market liquidity and transaction costs, in order to optimise the overall risk-return ratio [12].

The dynamic adjustment and optimisation of risk management is one of the key benefits of deep reinforcement learning. In traditional methods, risk management strategies are usually based on historical data and static rules, making it difficult to respond to real-time changes in the market environment. Deep reinforcement learning, on the other hand, enables risk management strategies to be continuously optimised in response to market changes through online learning and model updating. For example, in the event of sudden and dramatic market fluctuations, deep reinforcement learning models can quickly identify risk signals and instantly adjust strategies to avoid or mitigate losses from market shocks. In addition, deep reinforcement learning can automatically adjust the parameters of risk management strategies under different market conditions to adapt to different market volatility and trading environments. This dynamic adjustment capability makes deep reinforcement learning a significant advantage for risk management in high-frequency trading.

In conclusion, the application of deep reinforcement learning for risk management in high-frequency trading provides a new perspective and tool for traditional methods. Through dynamic adjustment and real-time optimisation, deep reinforcement learning can not only improve the flexibility and adaptability of trading strategies, but also control risks more effectively and safeguard the stability and continuity of high-frequency trading. In the future, with the further development of technology, the application of deep reinforcement learning in risk management is expected to become the core competitiveness in the field of high-frequency trading.

6. Discussion

The application of Deep Reinforcement Learning (DRL) in high-frequency trading shows great promise, but current research and application still face many challenges and limitations. Future research should focus on overcoming these obstacles and further deepening the integration of DRL with high-frequency trading to achieve higher efficiency, stability, and risk control.

At this stage, the application of DRL in high-frequency trading faces the contradiction between computational resources and real-time demand, as well as the possible strategy failure problem when dealing with non-stationary market environments. High-frequency trading requires decisions to be made in a very short period of time, and DRL algorithms often require a large amount of computational resources and training time, which makes it difficult for them to realise their full potential in real trading. In addition, due to the rapid changes in the market environment, DRL models may not be able to adapt in a timely manner when dealing with non-stationary data, leading to a reduction in the effectiveness of the strategies in real-world applications. To cope with these issues, future research should focus on improving the robustness and generalisation of the algorithms, especially in dealing with complex market environments.

Further integration of DRL with high-frequency trading in future research will focus on the following aspects: first, optimising algorithms and computational power to improve the real-time and adaptability of DRL in high-frequency trading. By introducing more complex models and optimisation strategies, DRL can better adapt to dynamic changes in the market, adjust trading strategies in real time, and maximise returns while effectively controlling risks. In addition, the introduction of cutting-edge technologies such as quantum computing is expected to significantly improve DRL's ability to handle ultra-high-dimensional data and complex environments, breaking through the current computational bottleneck.

Secondly, future research should also devote more attention in the direction of optimising risk management strategies. Current DRL models have the challenge of balancing return and risk in risk management, especially when facing extreme market conditions, which may not allow them to make optimal decisions in a timely manner. Future research can explore a more intelligent and dynamic risk control mechanism, adaptively adjusting risk exposure through DRL models and dynamically adjusting risk management strategies based on real-time market data. At the same time, the DRL model's ability to predict and respond to risks under extreme market conditions can be enhanced by combining scenario analyses and stress tests to improve the robustness of the overall system.

In addition, the application of multi-intelligence system (MAS) and integrated learning will also play an important role in future research on high-frequency trading and risk management. In complex market environments, it is difficult for a single intelligence to fully capture market dynamics. Through MAS, multiple intelligences can work together to achieve more refined and comprehensive market decisions. Integrated learning, on the other hand, can improve the accuracy and stability of trading decisions and reduce model bias and variance by combining the prediction results of multiple models.

Overall, future research should delve deeper into improving the robustness of DRL algorithms, optimising risk management strategies, introducing multi-intelligent body systems and integrated learning, and incorporating emerging technologies. These efforts will promote the application of DRL in high-frequency trading towards higher intelligence and efficiency, and provide strong technical support for the stable and efficient operation of financial markets.

7. Conclusion

This paper systematically explores the problem of optimisation and risk management of high-frequency trading algorithms based on Deep Reinforcement Learning (DRL) and provides a comprehensive assessment of existing research. By reviewing and analysing current research results, we can see that the application of Deep Reinforcement Learning in high-frequency trading shows remarkable potential. Its unique adaptive capabilities and dynamic decision-making mechanisms enable it to effectively optimise trading strategies and manage risks in complex and rapidly changing market environments. However, although DRL has made some progress in this field, it still faces many challenges, such as the real-time nature of the algorithms, data sparsity, overfitting problems and the complexity of risk management, which need to be further addressed in future research.

In the course of our research, we found that deep reinforcement learning can effectively capture potential patterns in the market through its powerful learning ability and achieve better revenue returns in high-frequency trading. Meanwhile, the application of DRL in risk management provides the possibility of dynamic adjustment for traditional static risk control methods, which significantly improves the robustness of the system and its ability to cope with unexpected market events. However, DRL models are prone to overfitting when dealing with high-dimensional market data, and the performance of the models may be significantly degraded especially under extreme market conditions. In addition, the performance of DRL in real market environments is limited by the complexity of the simulation environment, and current training environments are not yet able to fully reproduce all the dynamics in the market. The existence of these problems suggests that although DRL brings new methods and tools for high-frequency trading, in practice, it still needs to be further optimised and adapted to achieve the best results.

Based on the analyses in this paper, we make the following recommendations for future research. First, research should continue to explore how to enhance the real-time and robustness of DRL models by improving the algorithmic structure and optimising the training process to better suit the needs of high-frequency trading. Second, how to effectively deal with the sparsity and complexity of market data to avoid the overfitting problem of the model is a key direction for future research. Research can try to introduce more complex regularisation techniques and data enhancement methods to improve the generalisation ability of the model. In addition, optimisation of risk management strategies remains an important research area. Future research can combine DRL with multi-intelligence systems to develop more intelligent and flexible risk management frameworks by integrating learning so as to achieve better risk control in dynamic markets.

In conclusion, the application of deep reinforcement learning in high-frequency trading algorithm optimisation and risk management is promising, but it still needs to overcome the existing technical challenges through continuous research and practice. With the continuous evolution of algorithmic technology and market environment, DRL is expected to play a more important role in the future financial market and promote the further development and improvement of high-frequency trading technology.

References

[1] Li, J., Zhang, Y., Yang, X., & Chen, L. (2023). Online portfolio management via deep reinforcement learning with high-frequency data. Information Processing & Management, 60(3), 103247.

[2] Shavandi, A., & Khedmati, M. (2022). A multi-agent deep reinforcement learning framework for algorithmic trading in financial markets. Expert Systems with Applications, 208, 118124.

[3] Hirchoua, B., Ouhbi, B., & Frikh, B. (2021). Deep reinforcement learning based trading agents: Risk curiosity driven learning for financial rules-based policy. Expert Systems with Applications, 170, 114553.

[4] Rundo, F. (2019). Deep LSTM with reinforcement learning layer for financial trend prediction in FX high frequency trading systems. Applied Sciences, 9(20), 4460.

[5] Sahu, S. K., Mokhade, A., & Bokde, N. D. (2023). An overview of machine learning, deep learning, and reinforcement learning-based techniques in quantitative finance: recent progress and challenges. Applied Sciences, 13(3), 1956.

[6] Lin, Y. C., Chen, C. T., Sang, C. Y., & Huang, S. H. (2022). Multiagent-based deep reinforcement learning for risk-shifting portfolio management. Applied Soft Computing, 123, 108894.

[7] Wang, R., Wei, H., An, B., Feng, Z., & Yao, J. (2020). Deep stock trading: A hierarchical reinforcement learning framework for portfolio optimization and order execution. arXiv preprint arXiv:2012.12620.

[8] Han, L., Ding, N., Wang, G., Cheng, D., & Liang, Y. (2023, August). Efficient Continuous Space Policy Optimization for High-frequency Trading. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 4112-4122).

[9] Salehpour, A., & Samadzamini, K. (2023). Machine learning applications in algorithmic trading: a comprehensive systematic review. International Journal of Education and Management Engineering, 13(6), 41.

[10] Sun, S., Xue, W., Wang, R., He, X., Zhu, J., Li, J., & An, B. (2022, October). DeepScalper: A risk-aware reinforcement learning framework to capture fleeting intraday trading opportunities. In Proceedings of the 31st ACM International Conference on Information & Knowledge Management (pp. 1858-1867).

[11] Lei, K., Zhang, B., Li, Y., Yang, M., & Shen, Y. (2020). Time-driven feature-aware jointly deep reinforcement learning for financial signal representation and algorithmic trading. Expert Systems with Applications, 140, 112872.

[12] Millea, A., & Edalat, A. (2022). Using deep reinforcement learning with hierarchical risk parity for portfolio optimization. International Journal of Financial Studies, 11(1), 10.