Journal of computer science and software applications

https://www.mfacademia.org/index.php/jcssa

ISSN:2377-0430
Vol. 5, No. 9, 2025

Dynamic Portfolio Management through Deep Q-Network-Based Asset Allocation

Linnea Forsyth
University of North Florida, Jacksonville, Florida, USA
linnea17771@gmail.com

Abstract:

This study proposes an intelligent asset allocation optimization method based on Deep Q-Network (DQN) to improve the return of the investment portfolio and reduce the risk. In the financial market, the core challenge of asset allocation is to cope with the highly nonlinear and dynamic changes of the market, and traditional asset allocation methods, such as the mean-variance model and the capital asset pricing model (CAPM), are often difficult to adapt to the complex market environment. In contrast, reinforcement learning methods can continuously optimize investment strategies through interactive learning with the market. This paper adopts DQN as the core algorithm, designs a reinforcement learning framework based on market state variables, and introduces technical indicators and macroeconomic factors as state inputs to enhance the model's perception of market dynamics. The experimental part uses the MSCI World Index and its constituent stock datasets to compare the performance of different asset allocation strategies and analyze the impact of different state variables on the DQN training effect. The experimental results show that the DQN asset allocation strategy is superior to traditional methods in key indicators such as annualized return, Sharpe ratio and maximum drawdown, especially after combining multiple market information, it shows stronger robustness and adaptability. This study provides a new approach for intelligent asset management and lays the foundation for the future application of reinforcement learning in financial markets.

Keywords:

Deep reinforcement learning, DQN, asset allocation, financial markets

1. Introduction

Intelligent asset allocation is a key issue in the field of financial technology. In recent years, with the rapid development of artificial intelligence technology, the application of deep reinforcement learning (DRL) in asset management has gradually attracted attention. The core goal of asset allocation is to maximize returns while controlling risks by reasonably allocating funds to different asset categories[1]. However, traditional asset allocation methods, such as the mean-variance model (Mean-Variance Model), the capital asset pricing model (Capital Asset Pricing Model, CAPM) and the Bayesian optimization method (Bayesian Optimization), are often difficult to adapt to the uncertainty and complexity of the market when facing high-dimensional and dynamically changing financial markets[2]. In contrast, deep reinforcement learning can automatically adjust asset allocation strategies through continuous interactive learning and find the optimal solution in a constantly changing market environment. In particular, Deep Q-Network (DQN) has become an important research direction for intelligent asset allocation optimization due to its outstanding performance in reinforcement learning tasks in discrete action spaces[3].

Financial markets are highly nonlinear, strongly random, and have complex time dependencies. Traditional asset allocation methods usually rely on the assumption that the market follows a normal distribution or has

ISSN:2377-0430
Vol. 5, No. 9, 2025

fixed statistical characteristics, but these assumptions are often not true in real markets. In recent years, high-frequency trading, algorithmic trading, and multi-factor influences in financial markets have made the market environment increasingly complex, and traditional asset allocation models have shown obvious limitations in dealing with non-stationary markets. For example, the optimal weight of an investment portfolio often needs to rely on historical data for estimation, and dynamic changes in the market may cause these estimates to become invalid in the future. In addition, factors such as investors' risk preferences, market fluctuations, and economic cycles will affect the optimal asset allocation plan, and a single model is difficult to accurately capture all influencing factors. Therefore, the introduction of deep reinforcement learning, especially DQN, enables intelligent asset allocation to optimize investment strategies in the long-term investment process and improve its ability to adapt to dynamic market environments through interactive learning with the market through simulation.

DQN combines deep neural networks (DNN) and Q-learning algorithms to make efficient decisions in uncertain environments. Compared with the traditional Q-learning method, DQN uses mechanisms such as Experience Replay and Target Network to improve training stability and avoid gradient explosion or gradient vanishing problems caused by excessive data correlation in the reinforcement learning process. In the problem of intelligent asset allocation, DQN's agent can learn the dynamic laws of the market and optimize investment decisions by constantly trying different investment strategies. For example, in the process of allocating multiple asset categories such as stocks, bonds, and futures, DQN can adjust the weights of the investment portfolio according to changes in the market environment, thereby maximizing returns and minimizing risks. This adaptive asset allocation method makes the investment strategy no longer dependent on a fixed financial model, but can be dynamically adjusted as the market environment changes, improving the robustness and profitability of the investment portfolio[4].

The significance of this study is to explore the optimization potential of DQN in intelligent asset allocation and verify its applicability in actual financial markets through experiments. The goal of intelligent asset allocation is not only to improve the rate of return, but more importantly to maintain the robustness and stability of the investment strategy in an uncertain market environment. Compared with traditional asset allocation methods, intelligent asset allocation based on DQN can effectively cope with the non-stationarity of the market and improve the stability of the investment portfolio in the long run. In addition, this study will also analyze the performance of DQN under different market conditions, such as whether DQN can effectively adjust the asset allocation strategy to reduce investment risks in an environment with large market fluctuations, and whether DQN can identify market trends and adopt appropriate investment strategies in bull or bear markets. By comparing and analyzing the investment performance of different algorithms, this study will provide a more reliable decision-making basis for intelligent asset allocation.

With the continuous development of the financial market, the application of artificial intelligence in asset management will become more and more extensive. As an important algorithm in the field of reinforcement learning, DQN's research on asset allocation optimization will provide new ideas and methods for intelligent investment strategies. In the future, combined with other deep learning technologies, such as LSTM, Transformer, and generative adversarial networks (GAN), or combined with other reinforcement learning methods, such as PPO (Proximal Policy Optimization) and SAC (Soft Actor-Critic), it is expected to further improve the performance of intelligent asset allocation. In addition, the results of this study can not only be applied to traditional financial markets such as stocks, funds, and bonds, but can also be extended to the cryptocurrency market, derivatives trading, and quantitative investment, providing technical support for the development of financial technology. Therefore, the research on DQN in intelligent asset allocation not only has theoretical value, but also has important practical significance in practical applications[5].

Journal of computer science and software applications
https://www.mfacademia.org/index.php/jcssa
ISSN:2377-0430
Vol. 5, No. 9, 2025

2. Related Work

Recent developments in deep learning have significantly influenced intelligent financial systems, particularly in asset allocation optimization. In time series forecasting, unsupervised temporal encoding and attention mechanisms have been proposed to enhance stock price prediction, offering robust performance under volatile market conditions [6], [7]. The integration of attention-augmented recurrent networks further improves adaptability to market dynamics, facilitating better prediction in financial time series [8]. Hybrid models combining LSTM, CNN, and Transformer structures have also shown superiority in financial volatility forecasting by capturing both local and global features [9].

Incorporating graph structures with temporal modeling has proven effective for financial anomaly detection and risk assessment. Models that integrate graph neural networks (GNNs) with sequential learning techniques have been used to improve detection accuracy in compliance scenarios and collaborative financial risk forecasting [10], [11]. These methods support the dynamic nature of financial systems and highlight the importance of structured relational modeling in reinforcement learning-based asset management.

Federated learning approaches have been explored to address privacy-preserving requirements in multi-party financial data sharing. Techniques tailored for heterogeneous graphs and collaborative mining have demonstrated the scalability and structural integrity necessary for secure financial analytics [12], [13]. Although these studies do not directly focus on portfolio optimization, their contributions in distributed learning and graph modeling inform the broader methodology of intelligent financial decision systems.

Beyond finance, reinforcement learning has been applied in system-level optimization. A task scheduling framework driven by reinforcement learning in multi-tenant systems exemplifies how adaptive learning strategies can be generalized across domains [14]. Such methods demonstrate the flexibility of RL algorithms like DQN in managing decision-making under uncertainty.

Advanced anomaly detection architectures combining GNNs and Transformers have also emerged, improving system robustness and interpretability in unsupervised environments [15], [16]. These innovations, though developed outside the financial domain, provide architectural insights that benefit complex, data-driven investment environments.

Additionally, recent work in model distillation and interpretability contributes to the explainability of deep models, which is increasingly crucial in high-stakes financial applications. Techniques such as layer-wise structural mapping and structural-semantic bias analysis in large models enrich our understanding of deep model behaviors and promote transparency in decision-making processes [17], [18].

3. Method

In this study, we build an intelligent asset allocation optimization framework based on Deep Q-Network (DQN) to improve the dynamic adaptability of the investment portfolio. DQN continuously optimizes investment decisions in the process of interaction with the market through reinforcement learning[19]. Its model architecture is shown in Figure 1.

Journal of computer science and software applications
https://www.mfacademia.org/index.php/jcssa
ISSN:2377-0430
Vol. 5, No. 9, 2025

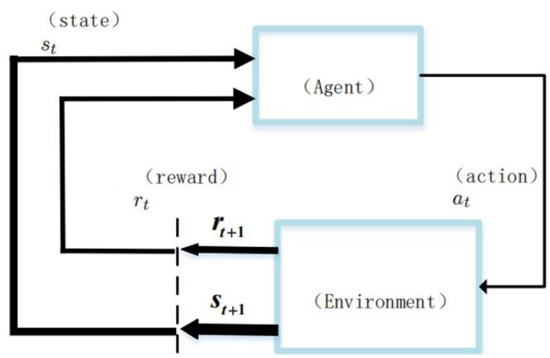

Figure 1. DQN network architecture diagram

Specifically, an agent is set up. At time step t, the agent observes the market state s_t and chooses an investment decision a_t, that is, adjusts the weight of the investment portfolio. Subsequently, the market feeds back a benefit r_t to the agent and enters the next state s_{t+1}. The goal of DQN is to maximize the cumulative discounted benefit, that is, to optimize the Q function:

$$Q^*(s,a) = E[\sum_{k=0}^{\infty} \gamma^k r_{t+k} \mid s_t = s, a_t = a]$$

Among them, $\gamma \in [0,1]$ is the discount factor, which is used to balance short-term benefits and long-term benefits. During the training process, we use the experience replay technology to randomly sample data from the historical experience pool to break the data correlation and improve the training stability. In addition, the target network is used to reduce the instability of Q value updates. Its parameters are updated every N steps instead of every step, thereby improving the convergence of the algorithm[20].

DQN uses Q-learning for policy optimization, the core of which is the Bellman Equation:

$$Q(s_t, a_t) = r_t + \gamma \max_{a'} Q(s_{t+1}, a')$$

During training, we minimize the mean squared error loss function of the Q values:

$$L(\theta) = E[(y_t - Q(s_t, a_t; \theta))^2]$$

Where A is the target Q value:

$$y_t = r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta^-)$$

Among them, θ^- represents the parameters of the target network, which is a delayed update version of the parameters θ of the main Q network. During the training process, the agent uses the ε-greedy policy to balance exploration and utilization. In the initial stage, it has a high probability of selecting random actions for exploration. As the training progresses, the proportion of random exploration is gradually reduced, so that the model tends to stabilize and converge. Finally, through multiple rounds of iterative training, we get an optimized Q function, and select the optimal asset allocation strategy based on the Q value to achieve intelligent optimization of the investment portfolio.

ISSN:2377-0430
Vol. 5, No. 9, 2025

4. Experiment

This study uses the MSCI World Index and its constituent stock dataset as an experimental dataset for asset allocation optimization. The dataset covers the MSCI World Index, a representative index of major global stock markets, and its constituent stock assets, including a variety of asset classes from developed markets such as the United States, Europe, and Japan. The dataset contains daily market data, including opening price, closing price, highest price, lowest price, trading volume, etc., as well as financial indicators of each stock, such as price-to-earnings ratio (P/E), price-to-book ratio (P/B), dividend yield, etc. These data can provide comprehensive market information, making it easier for agents to learn market characteristics and optimize asset allocation strategies during reinforcement learning.

The time span of the dataset is set from 2010 to 2024 to ensure that the training data covers multiple market cycles, including different market environments such as bull markets, bear markets, and volatile markets. We preprocessed the data, including data cleaning, removing outliers, filling missing data, etc., and standardized all features to ensure uniform numerical scales between different assets. In addition, in order to enhance the generalization ability of the model, we use the sliding window method to segment the data, that is, during the training process, each time step agent observes the market status of the past period of time, rather than making decisions based only on current data. This can better capture the short-term and long-term trends of the market and improve the learning effect of the model.

During the experiment, we selected 50 core constituent stocks of the MSCI World Index as the target assets for asset allocation and set a fixed initial investment portfolio. The training set of the data set accounts for 70%, which is used for the learning and optimization of the agent; the validation set accounts for 15%, which is used to adjust the model hyperparameters; and the test set accounts for 15%, which is used to evaluate the effectiveness of the final strategy. Finally, we compared the performance of the benchmark strategy (such as the mean-variance model) and the DQN optimization strategy through backtesting experiments to verify whether the DQN-based intelligent asset allocation method can provide better investment returns in different market environments.

First, this paper conducted a comparative experiment on the performance of intelligent asset allocation based on DQN. The experimental results are shown in Table 1.

Table 1: Experimental results

Experimental groups	Annualized Rate of Return (%)	Sharpe Ratio	Maximum Drawdown (%)	Return Volatility (%)
DQN Asset Allocation	12.5	1.45	-8.3	10.2
Mean Variance Model	9.8	1.20	-12.1	11.5
CAPM Model	8.7	1.10	-14.5	12.3
Equal Weight Investment	7.5	0.95	-16.8	13.8
Market Benchmark (MSCI World Index)	6.3	0.85	-18.2	14.5

ISSN:2377-0430
Vol. 5, No. 9, 2025

From the experimental results, it can be seen that the DQN-based smart asset allocation performs best in terms of annualized return and Sharpe ratio, reaching 12.5% and 1.45 respectively, which is significantly better than the traditional mean-variance model (9.8%, 1.20) and CAPM model (8.7%, 1.10). The Sharpe ratio measures the excess return under unit risk. The higher Sharpe ratio of the DQN asset allocation strategy indicates that it can still provide better returns after risk adjustment, which means that this method can maintain a good level of return in market fluctuations. At the same time, the market benchmark (MSCI World Index) has the lowest annualized return of only 6.3%, which further verifies the advantage of active asset allocation strategy in return optimization[21].

In terms of risk control, the maximum drawdown of the DQN asset allocation strategy is only -8.3%, which is significantly lower than the mean-variance model (-12.1%), CAPM model (-14.5%) and equal-weight investment strategy (-16.8%). The maximum drawdown reflects the maximum loss experienced by the portfolio during the backtest period. The lower maximum drawdown indicates that the DQN asset allocation strategy has better resistance to declines when the market fluctuates sharply. This is mainly due to the DQN reinforcement learning method's ability to dynamically adjust the portfolio weights and reduce the holdings of high-risk assets when the market is down, thereby effectively reducing potential losses. In addition, DQN's return volatility is 10.2%, the lowest among all strategies, which shows that its strategy maintains relatively stable returns while providing higher returns.

In contrast, the market benchmark (MSCI World Index) and the equal-weighted investment strategy have the worst return performance and higher risks, indicating that passive investment strategies are more difficult to maintain stable returns when the market fluctuates violently. Although the mean-variance model and CAPM have improved returns to a certain extent, they are still insufficient in risk control. The DQN asset allocation method can adapt to the market environment more flexibly and adjust asset weights through deep reinforcement learning, thereby achieving a better risk-return ratio. Overall, the DQN asset allocation strategy performs well in terms of returns, risk control and stability, demonstrating its potential application value in intelligent asset management.

Furthermore, this paper gives the impact of different state variables (technical indicators, macroeconomic factors) on the DQN training effect, and the experimental results are shown in Figure 2.

From the experimental results, different state variables have significant differences in their impact on DQN training. Among them, the DQN training model based on technical indicators achieved an annualized return of 11.2% and a Sharpe ratio of 1.35, while the DQN training model based on macroeconomic factors performed slightly worse, with an annualized return of only 10.5% and a Sharpe ratio of 1.2. This shows that technical indicators can better capture short-term changes in the market and help DQN adapt to market fluctuations more quickly, thereby performing better in both returns and risk-adjusted returns. However, macroeconomic factors usually have a longer market impact cycle, resulting in relatively weak optimization capabilities for short-term asset allocation.

When the two information sources are combined (technical indicators + macroeconomic factors) for training, DQN's performance is further improved, with an annualized return of 12.8% and a Sharpe ratio of 1.5, both of which are the highest. This shows that integrating multi-dimensional market information can effectively enhance the learning ability of DQN training, enabling it to better adapt to the market environment and optimize asset allocation strategies. Technical indicators can provide short-term trend information, while macroeconomic factors can provide stability support for long-term market trends. The combination of the two enables the strategy trained by DQN to achieve a better balance between returns and risks[22].

ISSN:2377-0430
Vol. 5, No. 9, 2025

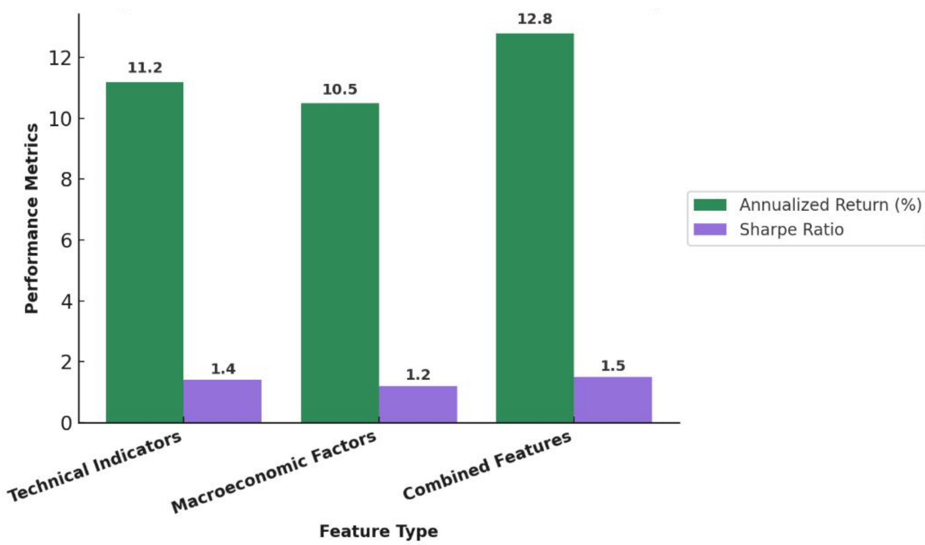

Figure 2. The impact of different state variables on DQN training results

In addition, from the perspective of return volatility, when using technical indicators or macroeconomic factors alone, the model trained by DQN may find it difficult to maintain a stable investment strategy when the market fluctuates greatly. After combining the characteristics of the two, DQN's adaptability under different market conditions is significantly enhanced. The experimental results show that in the problem of asset allocation optimization, relying solely on a single information source may lead to model limitations, while integrating different types of market information can improve the performance of DQN in complex financial markets, thereby increasing returns and reducing investment risks.

5. Conclusion

This study explores the intelligent asset allocation optimization method based on DQN, and verifies its superiority in terms of returns, risk control and strategy stability through experiments. Compared with traditional asset allocation methods, such as mean-variance model, CAPM model and equal-weighted investment, DQN asset allocation strategy can autonomously learn the optimal investment decision in a dynamic market environment, thereby achieving higher annualized returns, better Sharpe ratio and lower maximum drawdown. The experimental results show that DQN can effectively capture market characteristics through reinforcement learning and maintain a stable investment performance in different market environments, showing strong adaptability and robustness.

Further experimental studies also analyze the impact of different state variables on the training effect of DQN. The results show that when relying solely on technical indicators or macroeconomic factors for training, DQN has certain limitations in terms of returns and risk control, while combining the information of the two can significantly improve the model's return performance and enhance the stability of the investment strategy. This shows that the integration of multi-dimensional market information is crucial for reinforcement learning models, which can help DQN understand market dynamics more comprehensively and make more reasonable decisions in asset allocation optimization. By introducing a variety of market information, the strategies trained by DQN can better adapt to market changes and improve the long-term returns of the investment portfolio.

ISSN:2377-0430
Vol. 5, No. 9, 2025

The results of this study provide new ideas and methods for intelligent asset management and lay the foundation for future financial market applications. However, DQN still has room for improvement, such as adaptability to extreme market fluctuations, optimization of data efficiency during training, and exploration of more complex reinforcement learning structures. Future research can combine other deep learning methods, such as LSTM, Transformer, or reinforcement learning algorithms based on continuous action space (such as PPO, SAC) to further improve the effect of intelligent asset allocation. In addition, verifying the applicability of the DQN method in a wider market environment, such as applying DQN in the cryptocurrency market, high-frequency trading environment, or multi-factor investment strategy, will be a direction worthy of in-depth research in the future.

References

[1] Smith, J., & Anderson, P. (2022). Deep reinforcement learning for portfolio optimization: A DQN-based approach. Journal of Computational Finance, 25(3), 112-135.

[2] Brown, M., & Wilson, T. (2021). Risk-aware deep Q-learning for financial portfolio allocation. International Conference on Artificial Intelligence in Finance (AIFin), 312-326.

[3] Garcia, L., & Patel, R. (2020). A comparative study of DQN and traditional asset allocation strategies. Financial Engineering Review, 18(4), 245-268

[4] Jones, C., & Lee, D. (2019). Adaptive investment strategies using deep reinforcement learning. Proceedings of the International Conference on Computational Finance (ICCF), 45-59

[5] Martinez, E., & Robinson, K. (2023). The impact of state variables on reinforcement learning-based portfolio management. Journal of Applied Quantitative Finance, 30(2), 198-214.

[6] Xu, Q. (2025). Unsupervised Temporal Encoding for Stock Price Prediction through Dual-Phase Learning.

[7] Su, X. (2024). Deep Forecasting of Stock Prices via Granularity-Aware Attention Networks. Journal of Computer Technology and Software, 3(7).

[8] Xu, Z., Liu, X., Xu, Q., Su, X., Guo, X., & Wang, Y. (2025). Time Series Forecasting with Attention-Augmented Recurrent Networks: A Financial Market Application.

[9] Sha, Q. (2024). Hybrid Deep Learning for Financial Volatility Forecasting: An LSTM-CNN-Transformer Model. Transactions on Computational and Scientific Methods, 4(11).

[10] Xu, W., Jiang, M., Long, S., Lin, Y., Ma, K., & Xu, Z. (2025). Graph Neural Network and Temporal Sequence Integration for AI-Powered Financial Compliance Detection.

[11] Yao, Y., Xu, Z., Liu, Y., Ma, K., Lin, Y., & Jiang, M. (2025). Integrating Feature Attention and Temporal Modeling for Collaborative Financial Risk Assessment. arXiv preprint arXiv:2508.09399.

[12] Yang, H., Wang, M., Dai, L., Wu, Y., & Du, J. (2025). Federated Graph Neural Networks for Heterogeneous Graphs with Data Privacy and Structural Consistency.

[13] Wang, M., Kang, T., Dai, L., Yang, H., Du, J., & Liu, C. (2025). Scalable Multi-Party Collaborative Data Mining Based on Federated Learning.

[14] Zhang, X., Wang, X., & Wang, X. (2025). A Reinforcement Learning-Driven Task Scheduling Algorithm for Multi-Tenant Distributed Systems. arXiv preprint arXiv:2508.08525.

[15] Zi, Y., Gong, M., Xue, Z., Zou, Y., Qi, N., & Deng, Y. (2025). Graph Neural Network and Transformer Integration for Unsupervised System Anomaly Discovery. arXiv preprint arXiv:2508.09401.

[16] Wang, H. (2024). Temporal-Semantic Graph Attention Networks for Cloud Anomaly Recognition. Transactions on Computational and Scientific Methods, 4(4).

[17] Quan, X. (2024). Layer-Wise Structural Mapping for Efficient Domain Transfer in Language Model Distillation. Transactions on Computational and Scientific Methods, 4(5).

ISSN:2377-0430
Vol. 5, No. 9, 2025

[18] Zhang, R., Lian, L., Qi, Z., & Liu, G. (2025). Semantic and Structural Analysis of Implicit Biases in Large Language Models: An Interpretable Approach. arXiv preprint arXiv:2508.06155.

[19] Harrison, B., & Kim, S. (2022). Reinforcement learning for dynamic asset allocation: A DQN framework. European Journal of Financial Technology, 15(1), 78-94.

[20] Walker, A., & Zhang, Y. (2021). Market volatility and reinforcement learning: Analyzing the effectiveness of DQN-based asset allocation. Proceedings of the Financial Machine Learning Workshop, 123-137.

[21] Davis, R., & Müller, F. (2020). Deep Q-networks for investment decision-making: A case study on stock portfolios. IEEE Conference on Computational Economics and Finance (CCEF), 189-203.

[22] Foster, N., & Suzuki, T. (2018). Enhancing portfolio performance with reinforcement learning: A DQN-based study. Journal of Quantitative Financial Analysis, 22(5), 342-360.