International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

# Construction of Adaptive Portfolios using A Deep Reinforcement Learning for Risk Mitigation

***Rayavarapu Siri Pranathi [1]***

*B. Tech Student, Department of IT, GMR Institute of Technology, Rajam-532127, Andhra Pradesh, India*
*Email: 21341A12A6@gmrit.edu.in[1]*

**A B S T R A C T**

Choosing authentic stocks is critical to both quantitative trading and investing, which are intricate fiscal processes. The sensitive and big-time stochastic stock forecasting challenge has seen substantial evolution due to deep literacy advances, yet current results have two major downsides. They disregard the rich signals between associated stocks' temporal price oscillations and don't incontinently optimize the investment end in terms of profit. expanding upon these restrictions. Investors generally prefer portfolios with minimum correlations among their means in order to reduce the threat associated with connected means. The" Mercury" frame, a fresh system of stock webbing, is introduced in this approach. Your plutocrat grows as a effect of its capability to observe the request, acclimate to changes, and choose a collection of stocks that will not all vanish at formerly. Choosing the right stocks is pivotal to both quantitative trading and investing, which are intricate fiscal processes. The delicate and largely stochastic stock vaticination problem has seen substantial development due to deep literacy advances, yet current results have two major downsides. They disregard the rich signals between associated stocks' temporal price oscillations and don't incontinently optimize the investment end in terms of profit. expanding upon these restrictions. Investors generally prefer portfolios with minimum correlations among their means in order to reduce the threat associated with connected means. The" Mercury" frame, a new system of stock webbing, is introduced in this passage. Your plutocrat grows as a result of its capability to observe the request, acclimate to changes, and elect a collection of stocks that will not all vanish at formerly.

**Keywords:** ***Deep reinforcement learning, risk-return balanced portfolio strategy, market preferences, low-correlation assets.***

## Introduction

In contemporary investment strategies, constructing a stock portfolio serves dual purposes: mitigating investment risks through diversification and maximizing returns. Existing research predominantly employs machine learning and deep learning techniques to predict asset trends and select top-performing assets for portfolio optimization. However, many studies primarily focus on investment returns or incorporate macro-market data as risk indicators neglecting the potential benefits of qualitative information. While valuable historical data has limitations in capturing short-term returns, necessitating a more comprehensive approach. Qualitative information, such as news and events, plays a crucial role in decision-making, indicating the significance of incorporating financial text data into the investment process [11]. This integration enhances the understanding of the market and provides a robust foundation for making informed investment decisions, especially concerning long-term trends. Given the inherent risk in the stock market, investors seek a balanced portfolio that offers an optimal trade-off between risks and returns. Diversification across independent or minimally correlated asset classes is a common strategy to achieve this balance. The challenge lies in finding an effective strategy akin to reinforcement learning (RL), navigating the interplay between potential risks and rewards. In this work, we present a novel approach by formulating stock screening as a reinforcement learning process (Section III-A). The objective is to identify an optimal portfolio that strikes a balance between return and risk. Core modules integrated into the policy network include learning evolutionary features and inter-stock correlations, incorporating market sentiment indicators for long-term trend modeling, and evaluating individual stock performance comprehensively. The discrete and non-differentiable nature of market fluctuations and trading mechanisms is addressed through policy gradient optimization Extracting spatiotemporal features and inter-stock relations using hypergraph attention mechanisms, incorporating natural language processing sentiment indicators for long-term trend modeling, and evaluating stock performance at multi-granularity. Introduction of a novel ensemble framework, "Mercury," integrating reinforcement learning into the stock screening process to generate portfolios resilient to risk while ensuring returns. Seamless integration and joint training of all modules. Real-world experiments on

## Literature Review

Continuous Portfolio Optimization: The paper focuses on the problem of continuous portfolio optimization, which involves selecting the optimal allocation of assets from a continuous set, considering risk and return objectives.

Deep Reinforcement Learning (DRL): The authors propose a DRL-based approach to tackle the portfolio optimization problem. DRL has gained popularity in recent years for its ability to handle complex, high-dimensional tasks, making it suitable for portfolio optimization.

Adaptive Sampling Strategy: A key innovation in this work is the introduction of an adaptive sampling strategy. This strategy helps in efficient exploration of the continuous portfolio space, enabling the algorithm to discover optimal or near-optimal portfolios more effectively. Performance Evaluation: The paper provides extensive experiments and evaluations to demonstrate the effectiveness of the proposed algorithm. It compares the performance of the DRL-based approach with other traditional portfolio optimization methods, highlighting its advantages.

## Methodology

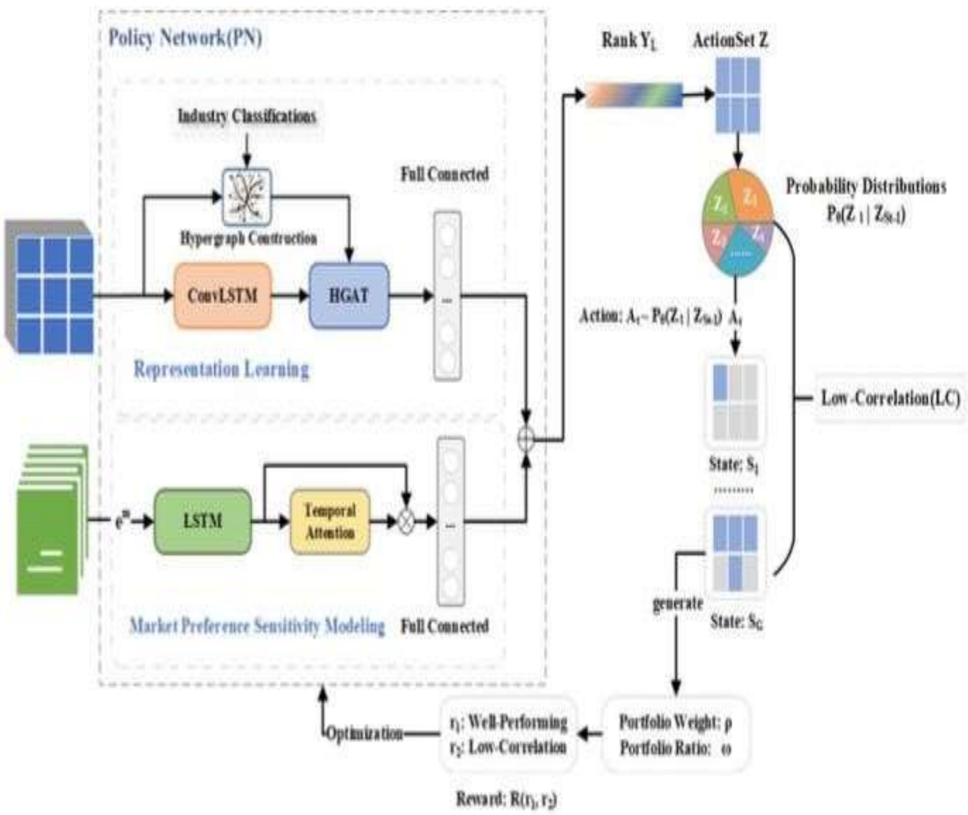


Strict stock selection is at the core of an intelligent investment decision-making process with the goal at developing diversified portfolios with little inter-asset correlation. Over time, the decision-making framework reveals an incremental method that includes foundational learning strategies for fully identifying low-correlated seeker stocks and optimising portfolios for returns that account for potential threats. The stock set called the Seeker

Point sequences (X_t) in (R{ T × D}) represent (S), encapsulating the numerous characteristics of every stock on a given trading day. Word vectors derived from stock bar commentary are used into request preference perceptivity modelling, whilst stock point sequences are used in representation literacy.

The covariance of (Z), which peaks in the firmness of state (SG) and names a low-correlation portfolio, determines the performance dimension. The set of actions The policy network (YL = PN(X, em)) is used to infer m, which provides direction for calculating portfolio weight. In order to balance the threat-return trade-off, the training optimization process uses policy-grade methods to align the performing portfolio with the requested criteria.

The framework moves through a number of interconnected steps, from the careful creation of a portfolio to the examination of stock input attributes and request emotion. underlying literacy strategies serve as crucial for choosing stocks with low correlation, which assists in accomplishing the overall goal of generating favorable returns while effectively avoiding hazards in the ever-changing landscape of investing making choices.

## Results

Mercury's performance was thoroughly assessed through a series of studies that shed light on a number of important areas of its functionality related to portfolio management. The tests utilised a dataset comprising fifty stocks from the Chinese A-share market, spanning the timeframe from June 2005 to December 2018.

Mercury performed better than RL-based techniques, conventional momentum strategies, and attention-based techniques, demonstrating its robustness and profitability. The model demonstrated long-term return robustness, particularly in difficult financial times like the subprime mortgage crisis. Mercury's endurance distinguished it from rivals such as Time Series Momentum and EIIE_LSTM, hence reinforcing its standing as a dependable instrument for portfolio optimisation.

The portfolio wealth on A-50.

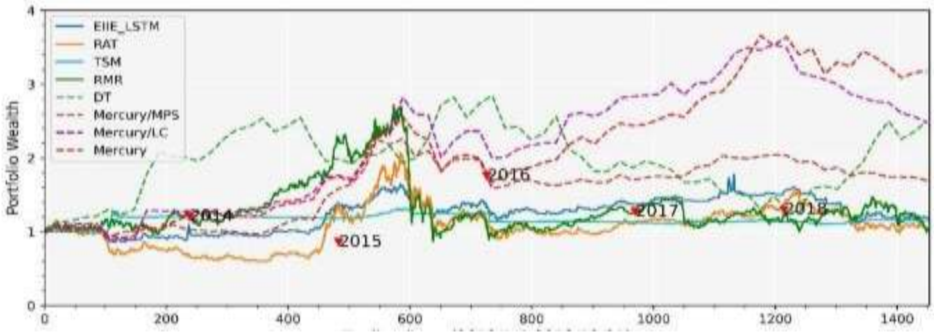


Mercury's fundamental modules were examined in more detail in ablation research, which highlighted the importance of Mercury/MPS in comprehending market risk. Mercury/LC responded similarly even though it did not have the low correlation (LC) structure, which highlighted the need of documenting individual stock performance regardless of correlation patterns.

TABLE 1. Performance comparisons on different models

| **Model** | **ARR** | **AVL** | **MDD** | **SR** | **CR** |
|---|---|---|---|---|---|
| **EIIE_LSTM** | 0.050 | 0.315 | 0.155 | 0.156 | 0.325 |
| **RAT** | 0.064 | 0.345 | 0.569 | 0.183 | 0.112 |
| **TSM** | 0.024 | 0.180 | 0.166 | 0.133 | 0.145 |
| **RMR** | 0.077 | 0.377 | 0.682 | 0.206 | 0.114 |
| **DT** | 0.091 | 0.275 | 0.463 | 0.292 | 0.174 |
| **Mercury/MPS** | 0.103 | **0.230** | **0.334** | 0.450 | 0.310 |
| **Mercury/LC** | 0.117 | 0.257 | 0.372 | 0.454 | 0.315 |
| **Mercury** | **0.125** | 0.245 | 0.338 | **0.741** | **0.538** |

Mercury's adaptability in addressing diverse investor preferences while upholding a strong risk tolerance was shown through the investigation of customised portfolios including distinct reward structures. Mercury's adaptability makes it a flexible option that can fit a wide variety of investment styles and tolerances for risk.

Understanding the trade-off among portfolios options and complexity was made possible through analysing the effects of various action set sizes and asset counts. The tests highlighted the need for a balanced approach, acknowledging that while additional assets offer potential diversity, they also raise complexity and transaction costs. Similarly, larger action sets offer more choices, but they also increase complexity. This complex insight adds to the model's flexibility in matching different investing objectives and approaches. To sum up, Mercury's full evaluation along four important axes validates its effectiveness, flexibility, and durability in the constantly evolving field of portfolio management. According to the findings, Mercury is a highly developed and reliable tool that can traverse challenging market conditions and perform very well in a variety of scenarios.

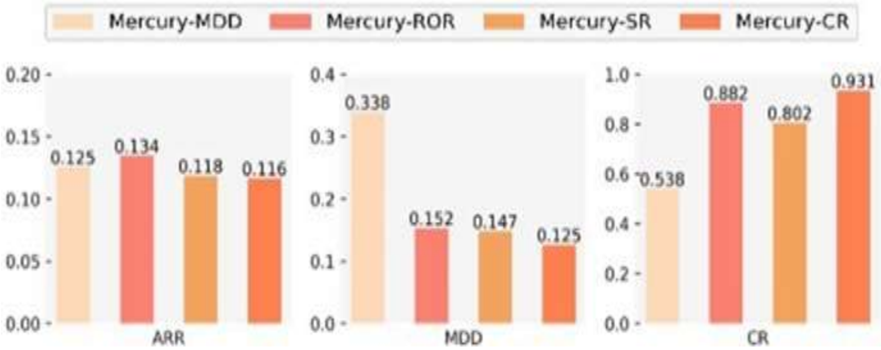


FIGURE : The performance of different reward functions.

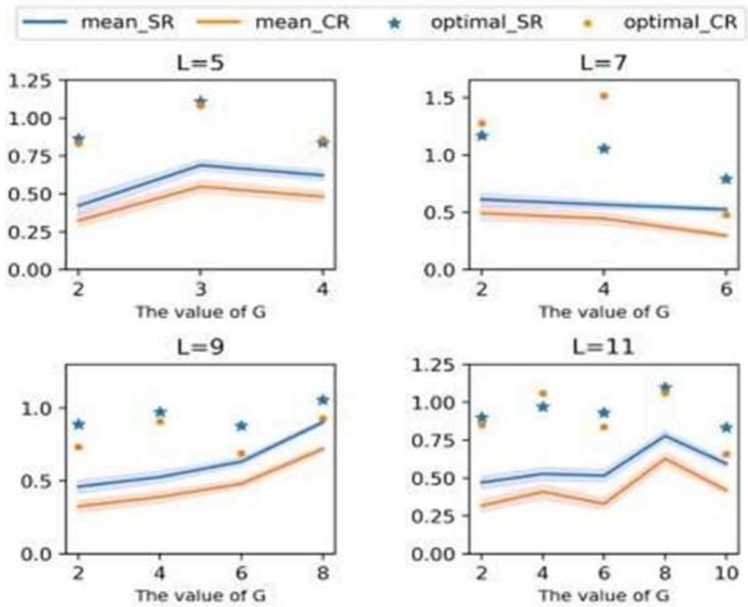


FIGURE : The performance of different values of L and G.

## Conclusion

Within the scope of reinforcement learning, we presented the Mercury. In order to identify assets that meet both tonalities, we merged historical data and stock market preferences in this work, which examined the relationships between stocks and market conditions. The reinforcement learning technique was extended to include stock screening, and market preferences were learned in order to jointly optimise the investing strategy with acceptable returns but low correlation between assets. In light of this, we integrated the fundamental modules into a policy network so that we could quickly swap them out for more sophisticated models in line with our next line of inquiry. To confirm the Mercury's efficacy, an actual A-share stock market experiment was conducted. And a straightforward adjustment to generalise a straightforward modification of the trading technique to the T + 0 market.

### References


1. R. Sawhney, S. Agarwal, A. Wadhwa, T. Derr, and R. R. Shah, “Stock Selection via Spatiotemporal Hypergraph Attention Network: A Learning to Rank Approach”, *AAAI*, vol. 35, no. 1, pp. 497-504, May 2021.
2. Z. Wang, B. Huang, S. Tu, K. Zhang, and L. Xu, “DeepTrader: A Deep Reinforcement Learning Approach for Risk-Return Balanced Portfolio Management with Market Conditions Embedding”, *AAAI*, vol. 35, no. 1, pp. 643-650, May 2021
3. S. -H. Huang, Y. -H. Miao and Y. -T. Hsiao, "Novel Deep Reinforcement Algorithm With Adaptive Sampling Strategy for Continuous Portfolio Optimization," in *IEEE Access*, vol. 9, pp. 77371-77385, 2021, doi: 10.1109/ACCESS.2021.3082186.
4. B. A. Luthfianti, D. Saepudin and A. F. Ihsan, "Portfolio Allocation of Stocks in Index LQ45 using Deep Reinforcement Learning," *2022 10th International Conference on Information and Communication Technology (ICoICT)*, Bandung, Indonesia, 2022, pp. 205-210, doi: 10.1109/ICoICT55009.2022.9914892.

5. N. Darapaneni *et al*., "Automated Portfolio Rebalancing using Q-learning," *2020 11th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, New York, NY, USA, 2020, pp. 0596-0602, doi: 10.1109/UEMCON51285.2020.9298035.
6. N. Darapaneni et al., "Automated Portfolio Rebalancing using Q-learning," 2020 11th IEEE Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON), New York, NY, USA, 2020, pp. 0596-0602, doi: 10.1109/UEMCON51285.2020.9298035.
7. L. Li, J. Wang and X. Li, "Efficiency Analysis of Machine Learning Intelligent Investment Based on K-Means Algorithm," in IEEE Access, vol. 8, pp. 147463-147470, 2020, doi: 10.1109/ACCESS.2020.3011366.
8. T. Kabbani and E. Duman, "Deep Reinforcement Learning Approach for Trading Automation in the Stock Market," in IEEE Access, vol. 10, pp. 93564-93574, 2022, doi: 10.1109/ACCESS.2022.3203697.
9. Y. -H. Miao, Y. -T. Hsiao and S. -H. Huang, "Portfolio Management based on Deep Reinforcement Learning with Adaptive Sampling," 2020 International Conference on Pervasive Artificial Intelligence (ICPAI), Taipei, Taiwan, 2020, pp. 130-133, doi: 10.1109/ICPAI51961.2020.00031.
10. U. Pigorsch and S. Schäfer, "High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning," 2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr), Helsinki, Finland, 2022, pp. 1-8, doi: 10.1109/CIFEr52523.2022.9776121.
11. B. Shi, X. Bai and C. Yao, "An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 11, pp. 2298-2304, 1 Nov. 2017, doi: 10.1109/TPAMI.2016.2646371.