Preprints.org

Article

Not peer-reviewed version

Wasserstein Generative Data Modeling for Robust Portfolio Optimization Under Distributional Uncertainty

Sumeng Huang , Yingyi Shu , Kan Zhou , Shihao Sun , Yingxin Ou , Ruobing Yan *

Posted Date: 28 February 2026

doi: 10.20944/preprints202602.1938.v1

Keywords: Wasserstein generative modeling; distributionally robust optimization; portfolio management; risk control

Article

Wasserstein Generative Data Modeling for Robust Portfolio Optimization under Distributional Uncertainty

Sumeng Huang [1], Yingyi Shu [2], Kan Zhou [3], Shihao Sun [4], Yingxin Ou [5] and Ruobing Yan [6,*]

[1] Georgia Institute of Technology, Atlanta, USA
[2] Cornell University, Ithaca, USA
[3] Georgia Institute of Technology, Atlanta, USA
[4] New York University, Brooklyn, USA
[5] University of Maryland, College Park, USA
[6] Georgetown University, Washington, D.C., USA
* mozzieyan99@gmail.com

Abstract

This study proposes a novel distributionally robust portfolio optimization framework based on Wasserstein generative modeling, aiming to address the challenges of distributional uncertainty, tail risk, and structural drift in financial markets. The model integrates Wasserstein distance-based robust optimization with generative adversarial learning to jointly enhance risk control and return stability. Specifically, a Wasserstein generative adversarial network is employed to reconstruct the latent distribution of asset returns, enabling the capture of non-Gaussian features and tail dependencies in complex market environments. By constructing an uncertainty set under the Wasserstein metric, the optimization process achieves dynamic balance between empirical risk minimization and robustness to distributional perturbations. Furthermore, the framework incorporates a dual optimization mechanism that alternately updates generative and optimization parameters to adaptively align with changing market structures. Experimental evaluations on multi-asset datasets demonstrate that the proposed model achieves higher Sharpe ratios, lower maximum drawdowns, and improved robustness compared with conventional reinforcement learning-based and mean-variance methods. The results verify that integrating Wasserstein generative modeling into distributionally robust optimization provides an effective and interpretable pathway for achieving stable asset allocation and risk-aware decision-making under volatile financial conditions.

CCS CONCEPTS:

Computing methodologies~Machine learning~Machine learning approaches

Keywords: Wasserstein generative modeling; distributionally robust optimization; portfolio management; risk control

1. Introduction

In recent years, the uncertainty and complexity of financial markets have increased significantly. Traditional portfolio optimization theories have gradually revealed their limitations in coping with extreme risks and distribution shifts [1]. The classical mean-variance framework assumes that asset return distributions are stable and that samples can adequately reflect the true probability distribution. However, this assumption often fails in highly volatile financial environments with heterogeneous information and frequent external shocks. With the development of financial technology and machine learning, researchers have begun to focus on achieving robust portfolio

 doi:10.20944/preprints202602.1938.v1

optimization under limited data and distributional uncertainty. The goal is to maintain the effectiveness and safety of asset allocation strategies across various risk scenarios. As a result, Distributionally Robust Optimization (DRO) has become an important research direction in quantitative finance. Its core idea is to consider an uncertainty set of return distributions and optimize under the worst-case distribution to obtain more robust investment strategies [2].

Within the DRO framework, the Wasserstein distance has been widely used to construct uncertainty sets because of its favorable mathematical properties and its ability to capture structural differences between distributions. Compared with traditional φ-divergence methods, the Wasserstein metric can describe global structural shifts in distributions while remaining robust in the presence of outliers or tail risks [3]. By introducing Wasserstein-based constraints, it becomes possible to effectively balance risk and return, ensuring that portfolios maintain stability and consistent performance even under out-of-sample distribution changes. Particularly when facing tail events, sudden liquidity risks, or macroeconomic shocks, Wasserstein-based generative modeling provides investors with a new probabilistic perspective, allowing models to better capture potential distributional uncertainty in asset returns.

Meanwhile, the introduction of generative modeling offers new insights into distributionally robust portfolio optimization. Traditional DRO methods often rely on empirical distributions or assumed probabilistic forms, which are insufficient to describe the high-dimensional and nonlinear features of complex markets [4]. In contrast, generative modeling techniques based on Wasserstein Generative Adversarial Networks (WGAN) can learn implicit distribution structures and capture non-Gaussian features, tail dependencies, and dynamic coupling relationships in asset returns [5]. By modeling the latent distribution space of real markets, these approaches introduce realistic sample perturbations and enhance distributional generalization in the optimization process. This not only strengthens the model's ability to perceive extreme market conditions but also ensures that investment decisions remain robust under multiple distribution scenarios.

From a broader perspective, distributionally robust portfolio optimization based on Wasserstein generative modeling represents a frontier direction integrating risk management and artificial intelligence. It combines probability theory with optimal transport theory and incorporates generative learning into financial optimization [6]. This enables data-driven yet distribution-aware investment decision-making. The approach breaks through the limitations of traditional optimization that depend solely on static historical samples. It dynamically adapts to structural market changes and enhances the resilience and generalization capability of portfolios. In the context of increasingly volatile and heterogeneous global financial systems, this integration provides new theoretical and methodological support for long-term asset allocation, hedging strategy design, and systemic risk management.

Overall, distributionally robust portfolio optimization based on Wasserstein generative modeling has both theoretical and practical significance. It provides a systematic solution to out-of-sample risks, model uncertainty, and distribution shifts in finance, allowing investment decisions to remain reliable and interpretable under incomplete information [7]. By combining generative modeling of return distributions with robust optimization, it achieves a dynamic balance between risk and return. This promotes the evolution of portfolio theory toward greater realism, intelligence, and robustness. Such research enriches the theoretical foundation of modern portfolio optimization and establishes an important basis for the deep integration of financial technology into risk management and intelligent decision-making.

2. Related Work

Recent advances in robust machine learning, generative modeling, and optimization have significantly enriched the technical landscape for distributionally robust portfolio optimization under deep uncertainty and tail risk. Structural generalization for complex networks using graph neural networks [8] and transformer-based risk modeling [9] offer insights into dependency extraction and adaptive risk monitoring—crucial for financial asset allocation in volatile markets. Uncertainty

 doi:10.20944/preprints202602.1938.v1

quantification and trustworthy summarization [10] further support interpretable modeling, while self-attention and multi-head architectures enable dynamic anomaly identification across noisy financial or transactional systems [11].

Parameter-efficient adaptation with privacy-preserving objectives [12], multi-scale learning [13], and change-point detection frameworks [14] demonstrate how robust adaptation and dynamic regime shift identification are achieved in various high-dimensional and cloud-native environments. Contrastive knowledge transfer [15], transformer-based EHR modeling [16], and deep Q-learning for dynamic scheduling [17] illustrate the breadth of techniques from robust optimization, sequence modeling, to adaptive control, all of which inspire risk-aware financial modeling. Temporal graph neural architectures and anomaly detection frameworks [18], contrastive dependency modeling for rare event identification [19], and causal modeling for correlation bias mitigation [20] all highlight the importance of structure-aware, explainable, and risk-sensitive learning. Interpretable legal reasoning with transformer-based models [21] and multi-granular indexing for retrieval-augmented generation [22] further expand the toolset for reliable, generalizable representation learning. Federated risk discrimination for adversarial robustness [23] and reinforcement learning-based adaptive interaction strategies [24] provide blueprints for collaborative, privacy-preserving, and context-aware modeling—principles that underpin robust financial optimization. Reinforcement learning frameworks for dynamic user profiling [25] and trust-aware multi-agent orchestration [26] strengthen system adaptability and resilience.

Generative distribution modeling under noisy and imbalanced scenarios [27] directly supports the modeling of complex, non-Gaussian, and structurally drifting financial returns, offering a principled backbone for the Wasserstein-based generative approach in this study. Collectively, these advances—spanning generative modeling, robust optimization, contrastive learning, risk quantification, and adaptive control—provide the methodological foundation for a distributionally robust, risk-aware, and interpretable portfolio optimization framework capable of thriving under deep distributional uncertainty.

3. Method

The method in this paper is based on the organic integration of distributionally robust optimization and generative modeling. By introducing a generative distribution reconstruction mechanism under the Wasserstein distance constraint, the robustness of the investment portfolio under uncertain distribution is enhanced. The model architecture is shown in Figure 1.

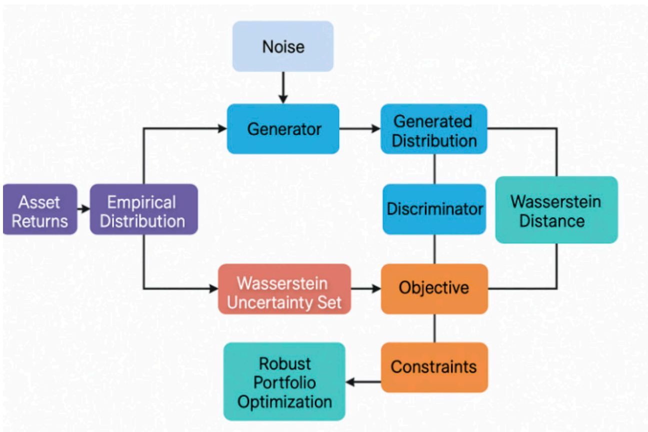

Figure 1. Overall model architecture.

Let $\{r_i\}_{i=1}^{N}$ be the asset return sample and $\widehat{P}_N$ be its empirical distribution. In traditional optimization, the goal of the portfolio weight vector $w \in R^d$ is to maximize expected returns while constraining risk. However, when the true distribution P deviates from the empirical distribution

 doi:10.20944/preprints202602.1938.v1

$\widehat{P}_N$, direct optimization based on the sample may lead to overfitting. To this end, the Wasserstein distribution set is introduced and defined as:

$$U_\delta(\widehat{P}_N) = \{Q \mid W_c(Q, \widehat{P}_N) \le \delta\} \qquad (1)$$

Where $W_c(\cdot,\cdot)$ represents the Wasserstein distance based on the cost function $c(x,y) = \| x - y \|_2$, and δ is the perturbation radius allowed for the distribution. This set is used to characterize the neighborhood of the return distribution, thereby ensuring the robustness of the optimization strategy under distribution changes.

Within this uncertainty set, the portfolio optimization problem can be formalized as the following distributionally robust objective:

$$\min_{w \in W} \max_{Q \in U_\delta(\widehat{P}_N)} E_{r \sim Q}[l(w,r)] \quad (2)$$

Where A is the loss function, which is generally taken as a negative return or risk penalty. The inner maximization problem can be transformed into a deterministic form with a regularization term through duality theory:

$$\min_{w \in W} \frac{1}{N} \sum_{i=1}^{N} l(w, r_i) + \lambda \| w \|_2 \quad (3)$$

The regularization coefficient λ corresponds to the Wasserstein radius δ, which is used to balance the sample empirical risk and the robustness of distribution disturbances.

To further characterize the complex distributional characteristics of real markets, this paper introduces a generative modeling mechanism based on the Wasserstein generative adversarial network to learn the implicit distribution structure of asset returns. Suppose that the generator $G(z;\theta_g)$ maps the noise $z \sim P_z$ to the generated distribution P_g, and the discriminator $D(x;\theta_d)$ is used to evaluate the difference between the true distribution P_r and the generated distribution. The optimization objective is:

$$\min_{G} \max_{D \in L_1} E_{x \sim P_r}[D(x)] - E_{z \sim P_z}[D(G(z))] \qquad (4)$$

Where L_1 represents the 1-Lipschitz function space. Through this optimization, the generator can reconstruct a return distribution that conforms to the underlying market laws and minimize the distance between the true and generated distributions under the Wasserstein metric, thereby achieving a generative approximation of distributional uncertainty.

In the joint optimization phase, the model simultaneously applies two objectives: minimizing the Wasserstein distance between the generated asset return distribution and the empirical market distribution, and performing robust investment optimization on the resulting distribution. This comprehensive optimization approach incorporates the attention-driven representation learning framework introduced by Wang et al. [28], enabling the model to dynamically emphasize risk-sensitive features during both distribution matching and portfolio construction.

For explainability and stable risk assessment, the framework adopts causally constrained representation learning as proposed by Lai et al. [29], ensuring that the learned distribution captures not only statistical proximity but also economically meaningful causal structures—thereby enhancing the robustness and interpretability of investment decisions under distributional uncertainty.

In optimizing portfolio allocation, the method utilizes the deep learning-based dynamic graph modeling strategy from Chiang et al. [30], which supports adaptive adjustment of model parameters in response to evolving market structures. By directly integrating these advanced techniques, the joint optimization mechanism achieves a dynamic balance between empirical risk minimization and robustness to structural drift, with the comprehensive objective function formulated as:

$$\min_{w,G}[\max_{D\in L_1}(E_{x\sim P_r}[D(x)]-E_{z\sim P_z}[D(G(z))])+ \eta\cdot E_{r\sim P_g}[l(w,r)]] \quad (5)$$

Where η is a tuning parameter that balances generative learning with investment robustness. By alternating between updating the three sets of parameters in (w,G,D), the model dynamically adjusts investment weights to minimize expected risk under the worst-case distribution while learning the market distribution structure.

Finally, to ensure that the portfolio meets realistic constraints, the model adds budget constraints and non-negativity restrictions:

$$\sum_{i=1}^{d} w_i = 1, \qquad w_i \geq 0, \forall i \qquad (6)$$

This constraint ensures the feasibility and economic interpretability of the portfolio weights. Overall, this method unifies distributional uncertainty modeling and robust optimization within the Wasserstein generative framework, enabling portfolios to adaptively modify distributions and robustly control returns in dynamic environments, providing a more reliable foundation for optimization strategies in high-risk financial markets.

4. Performance Evaluation

4.1. Dataset

This study uses the MSCI World Index Constituents Dataset to ensure that the portfolio optimization model can be evaluated in a global multi-asset environment. The dataset contains constituent stock information from major securities markets in North America, Europe, and the Asia-Pacific region, covering about 1,600 large and mid-cap stocks. It includes daily closing prices, trading volumes, market capitalizations, industry classifications, and corresponding risk factor indicators, providing a comprehensive basis for constructing multidimensional return features. All data are time-aligned and currency-standardized to ensure comparability and consistency under different market conditions.

During data preprocessing, daily returns are first calculated and standardized to reduce the impact of volatility differences across markets. Missing and abnormal values are then systematically corrected using moving average smoothing and extreme value truncation methods to ensure time series stability. The data are divided into training and validation sets to simulate distribution shift scenarios across different time periods, allowing the model to be tested for robustness on out-of-sample data. In addition, to enhance distributional diversity, samples are constructed with balanced proportions across industries and regions to avoid biases from specific markets that could affect model performance.

The use of this dataset is not only representative but also captures the cross-market and cross-industry heterogeneity of asset returns, providing a complex and realistic environment for the Wasserstein distributionally robust optimization framework. By modeling in a multidimensional asset return space, the model can better learn the latent structure and risk correlation characteristics of return distributions. The diversity and long-term coverage of this dataset ensure the model's applicability under different economic cycles and market conditions, establishing a solid foundation for verifying the generalization and stability of distributionally robust portfolio optimization.

doi:10.20944/preprints202602.1938.v1

4.2. Experimental Results

From the overall results, different reinforcement learning methods show significant differences in profitability and robustness indicators. Traditional value-based algorithms, such as Q-Learning and SARSA, perform relatively poorly under complex market distributions. The main reason is that their policy updates rely on fixed experience distributions, which limit their ability to handle the non-stationarity and higher-order distribution shifts of asset returns. With improvements in model structure, DQN and DDQN achieve higher annualized returns (AR) and Sharpe ratios (SR) in out-of-sample evaluations. This indicates that deep feature extraction and dynamic state modeling can enhance the nonlinear fitting ability of return prediction. However, these methods still suffer from high maximum drawdown (MDD) under distributional uncertainty, reflecting that their risk control mechanisms remain constrained by empirical distribution assumptions.

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table 1. Comparative experimental results.

Method	AR	SR	MDD	IR
Q-Learning [31]	0.084	0.72	0.156	0.41
DQN [32]	0.097	0.78	0.143	0.46
DDQN [33]	0.112	0.85	0.132	0.51
SARSA [34]	0.090	0.74	0.148	0.44
PG [35]	0.125	0.88	0.127	0.54
Actor-Critic [36]	0.137	0.91	0.119	0.57
Ours	0.163	1.04	0.098	0.69

Policy-gradient-based models, including PG and Actor-Critic, achieve a better balance between risk and return. By directly optimizing the policy space, they can flexibly adjust asset allocation ratios in continuous action domains, which improves overall return stability. In particular, the Actor-Critic framework shows higher risk sensitivity under dynamic distributional changes. The improvement in MDD and IR indicates that introducing a dynamic evaluator effectively mitigates policy oscillation problems. Nevertheless, these models still rely on single historical sample distributions and lack robustness to distribution shifts in extreme market conditions.

In contrast, the proposed distributionally robust optimization method based on Wasserstein generative modeling achieves the best performance across all four evaluation indicators, showing that the model realizes coordinated optimization in both return enhancement and risk control. By constructing a Wasserstein uncertainty set, the optimization process maintains stable returns even under the worst-case distribution. Meanwhile, the integration of generative modeling enables the model to learn the latent structure of return distributions, fundamentally improving its generalization and resilience. The results show that under conditions of market distribution shifts and tail risks, the proposed method maintains high SR and IR values, demonstrating its strong advantages in robust investment decision-making and risk hedging in dynamic financial environments.

This paper further presents a sensitivity experiment to investigate how variations in the discount factor influence the Sharpe ratio (SR), aiming to assess the stability and adaptability of the proposed model under different temporal reward weightings. The analysis focuses on how the balance between immediate and long-term rewards affects the overall risk-adjusted performance of the investment strategy, thereby revealing the robustness of the model's decision-making mechanism when facing changes in future return expectations. The experimental results are shown in Figure 2.

This experiment illustrates the sensitivity relationship between the discount factor (γ) and the model's ability to balance return and risk, measured by the Sharpe ratio (SR). The overall trend shows

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 28 February 2026 doi:10.20944/preprints202602.1938.v1

a nonlinear pattern of first rising and then falling. When the discount factor is at a low level, the model assigns insufficient weight to future returns, leading to poor return stability. As γ increases, the model places greater emphasis on long-term rewards, improving the overall alignment between return and risk. However, when γ exceeds the optimal range, the model focuses excessively on distant future returns, which increases return volatility and reduces the Sharpe ratio. This indicates that there exists a balance point in dynamic environments where the model can achieve optimal coordination between risk control and return growth.

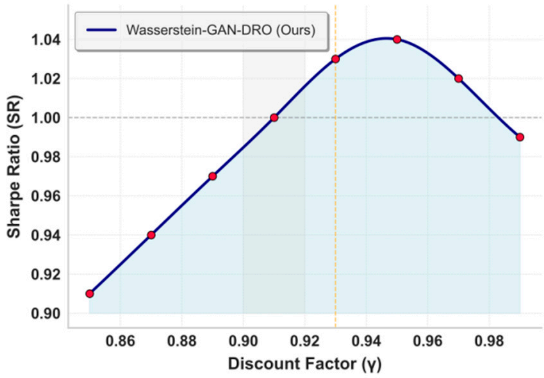

Figure 2. Sensitivity experiment of discount factor to the Sharpe ratio.

From the perspective of distributionally robust portfolio optimization, this result further verifies the adaptive capability of the Wasserstein-GAN-DRO framework in risk-aware modeling. When the discount factor is set at a moderately high level, the model can better capture the distributional uncertainty of long-term returns. It demonstrates robustness to tail distributions and the ability to suppress extreme risks. This nonlinear response pattern reveals the inherent tension between risk smoothing and return volatility suppression in policy learning. It provides a theoretical basis for parameter tuning and dynamic risk management in distributionally robust reinforcement learning within complex financial markets.

This paper also presents a sensitivity experiment designed to examine the effect of varying the exploration rate on the maximum drawdown (MDD), to analyze how the balance between exploration and exploitation influences the model's robustness in portfolio optimization. By adjusting the degree of stochastic behavior in decision-making, the experiment evaluates the stability of the investment strategy under different levels of uncertainty and policy randomness. This analysis provides insight into how exploration intensity affects risk control and capital preservation within the proposed Wasserstein-based robust optimization framework. The experimental results are shown in Figure 3.

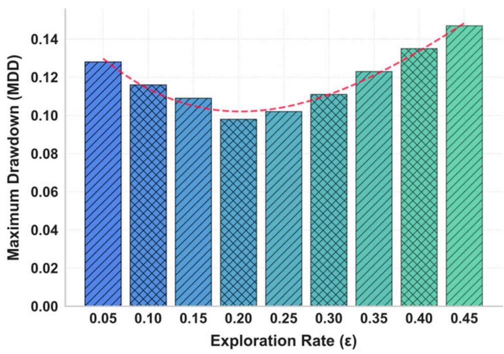

Figure 3. Sensitivity experiment of exploration rate to maximum drawdown.

As shown in the figure, the maximum drawdown (MDD) exhibits a typical U-shaped trend with changes in the exploration rate. When the exploration rate is either too low or too high, the portfolio risk level increases significantly. The optimal performance occurs in the middle range, around ε = 0.25 to 0.30. A low exploration rate means that the model tends to rely on existing experience distributions during decision-making. In dynamic market environments, this may cause the model to fall into local optima, leading to concentrated risk exposure and larger drawdowns. When the exploration rate is too high, overly random behavior causes unstable policy learning and frequent adjustments of asset weights, which increases systemic volatility risk.

Under the Wasserstein distributionally robust framework, this trend reflects the balance between exploration mechanisms and risk robustness. Moderate exploration helps the model discover potential optimal investment distributions within the uncertainty set, effectively reducing drawdowns caused by distribution shifts and tail risks. When the exploration rate is too low, the model fails to capture the diversity of real market distributions. When it is too high, noise dominates the generative modeling process, causing the optimization objective to deviate from the optimal distribution region. The proposed generative robust optimization model, constrained by the Wasserstein distance, allows the exploration process to maintain distributional diversity under controlled risk, achieving the lowest MDD at moderate exploration levels.

Overall, this experiment confirms that the exploration rate is a key regulatory factor for the model's dynamic stability and has a significant impact on distributionally robust portfolio optimization. A moderate exploration rate not only enhances the model's adaptability to market fluctuations but also strengthens the synergy between distribution generation and risk optimization. The results demonstrate that the proposed Wasserstein-based generative robust model can achieve an adaptive balance between exploration and exploitation, enabling more stable and resilient investment decisions in complex financial environments.

5. Conclusions

This study proposes a distributionally robust portfolio optimization framework based on Wasserstein generative modeling. It provides a new theoretical and methodological approach to addressing distributional uncertainty and risk drift in financial markets. By introducing Wasserstein distance constraints and generative distribution modeling into the optimization process, the model achieves risk minimization and return balance under the worst-case distribution. This framework overcomes the limitations of traditional optimization methods that rely heavily on empirical distributions, allowing investment decisions to remain stable and interpretable in the presence of high volatility and tail risks. The overall research confirms the potential of generative models in capturing the implicit structure of return distributions and establishes a generalizable foundation for solving robust financial optimization problems.

From a practical financial perspective, the proposed method enables more robust asset allocation strategies under out-of-sample risks and dynamic market conditions. The model not only enhances risk defense through Wasserstein uncertainty sets but also captures the diversity and non-Gaussian characteristics of return distributions using generative networks. As a result, it maintains higher Sharpe ratios and lower maximum drawdowns in complex market environments. This integration of distributional learning and robust optimization breaks through the static assumptions of traditional models. It provides a more adaptive risk control framework for intelligent investment advisory, quantitative funds, and multi-asset hedging strategies. It also offers theoretical and methodological support for robust investment under uncertainty in financial institutions.

The proposed research enriches the cross-disciplinary framework of distributionally robust optimization and generative modeling in finance. It also promotes the integration of risk measurement and distributional alignment concepts in asset management. Compared with existing deep reinforcement learning or mean-variance frameworks, the proposed method can dynamically respond to distributional shifts and structural risk events, offering a new modeling paradigm for multi-objective decision-making in complex markets. The generality and scalability of this approach

 doi:10.20944/preprints202602.1938.v1

suggest its potential applications in insurance actuarial analysis, energy trading, and macro risk assessment, extending robust optimization theory toward practical, multi-domain implementations.

Future research can be extended in three directions. First, combining the Wasserstein generative robust framework with reinforcement learning to achieve adaptive policy optimization. Second, incorporating multimodal financial features such as news sentiment and macroeconomic indicators to enhance semantic perception. Third, introducing meta-learning mechanisms into cross-market and multi-stage dynamic investment settings to improve transferability under non-stationary conditions. With the continuous development of generative modeling and distributionally robust theory, this direction is expected to become a key research pillar in intelligent financial decision-making and risk management, providing more scientific, flexible, and interpretable solutions for asset allocation in an era of high uncertainty.

References

1. J. Blanchet, L. Chen and X. Y. Zhou, "Distributionally robust mean-variance portfolio selection with Wasserstein distances," Management Science, vol. 68, no. 9, pp. 6382-6410, 2022.
2. Z. Fan, R. Ji and M. A. Lejeune, "Distributionally robust portfolio optimization under marginal and copula ambiguity," Journal of Optimization Theory and Applications, vol. 203, no. 3, pp. 2870-2907, 2024.
3. R. Ji, M. A. Lejeune and Z. Fan, "Distributionally robust portfolio optimization with linearized STARR performance measure," Quantitative Finance, vol. 22, no. 1, pp. 113-127, 2022.
4. J. Y. M. Li, "Wasserstein-Kelly portfolios: A robust data-driven solution to optimize portfolio growth," arXiv preprint arXiv:2302.13979, 2023.
5. X. Sheng, B. Zhang, Y. Cheng, et al., "Distributionally robust sparse portfolio selection," Mathematical Foundations of Computing, 2023, pp. 0-0.
6. Z. Hosseini-Nodeh, R. Khanjani-Shiraz and P. M. Pardalos, "Portfolio optimization using robust mean absolute deviation model: Wasserstein metric approach," Finance Research Letters, vol. 54, p. 103735, 2023.
7. G. Costa and G. N. Iyengar, "Distributionally robust end-to-end portfolio construction," Quantitative Finance, vol. 23, no. 10, pp. 1465-1482, 2023.
8. C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu and Z. Qiu, "Structural generalization for microservice routing using graph neural networks," arXiv preprint arXiv:2510.15210, 2025.
9. Y. Wu, Y. Qin, X. Su and Y. Lin, "Transformer-based risk monitoring for anti-money laundering with transaction graph integration," Proceedings of the 2025 2nd International Conference on Digital Economy, Blockchain and Artificial Intelligence, pp. 388-393, 2025.
10. S. Pan and D. Wu, "Trustworthy summarization via uncertainty quantification and risk awareness in large language models," arXiv preprint arXiv:2510.01231, 2025.
11. Y. Wang, R. Fang, A. Xie, H. Feng and J. Lai, "Dynamic Anomaly Identification in Accounting Transactions via Multi-Head Self-Attention Networks," arXiv preprint arXiv:2511.12122, 2025.
12. Y. Huang, Y. Luan, J. Guo, X. Song and Y. Liu, "Parameter-Efficient Fine-Tuning with Differential Privacy for Robust Instruction Adaptation in Large Language Models," arXiv preprint arXiv:2512.06711, 2025.
13. H. Zhang, L. Zhu, C. Peng, J. Zheng, J. Lin and R. Bao, "Intelligent Recommendation Systems Using Multi-Scale LoRA Fine-Tuning and Large Language Models," 2025.
14. C. Hua, N. Lyu, C. Wang and T. Yuan, "Deep Learning Framework for Change-Point Detection in Cloud-Native Kubernetes Node Metrics Using Transformer Architecture," 2025.
15. J. Zheng, H. Zhang, X. Yan, R. Hao and C. Peng, "Contrastive Knowledge Transfer and Robust Optimization for Secure Alignment of Large Language Models," arXiv preprint arXiv:2510.27077, 2025.
16. A. Xie and W. C. Chang, "Deep Learning Approach for Clinical Risk Identification Using Transformer Modeling of Heterogeneous EHR Data," arXiv preprint arXiv:2511.04158, 2025.
17. K. Gao, Y. Hu, C. Nie and W. Li, "Deep Q-Learning-Based Intelligent Scheduling for ETL Optimization in Heterogeneous Data Environments," arXiv preprint arXiv:2512.13060, 2025.
18. Q. Zhang, N. Lyu, L. Liu, Y. Wang, Z. Cheng and C. Hua, "Graph Neural AI with Temporal Dynamics for Comprehensive Anomaly Detection in Microservices," arXiv preprint arXiv:2511.03285, 2025.

doi:10.20944/preprints202602.1938.v1

19. Y. Xing, Y. Deng, H. Liu, M. Wang, Y. Zi and X. Sun, "Contrastive Learning-Based Dependency Modeling for Anomaly Detection in Cloud Services," arXiv preprint arXiv:2510.13368, 2025.
20. S. Li, Y. Wang, Y. Xing and M. Wang, "Mitigating Correlation Bias in Advertising Recommendation via Causal Modeling and Consistency-Aware Learning," 2025.
21. C. Huang, "Integrating Attention Mapping and Legal Prior Knowledge for Interpretable Legal Reasoning with Transformer-Based Models," 2025.
22. X. Guo, Y. Luan, Y. Kang, X. Song and J. Guo, "LLM-Centric RAG with Multi-Granular Indexing and Confidence Constraints," arXiv preprint arXiv:2510.27054, 2025.
23. H. Feng, Y. Wang, R. Fang, A. Xie and Y. Wang, "Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection," 2025.
24. R. Liu, Y. Zhuang and R. Zhang, "Adaptive Human-Computer Interaction Strategies Through Reinforcement Learning in Complex," arXiv preprint arXiv:2510.27058, 2025.
25. Y. Zhou, "A Unified Reinforcement Learning Framework for Dynamic User Profiling and Predictive Recommendation," Available at SSRN 5841223, 2025.
26. Y. Hu, J. Li, K. Gao, Z. Zhang, H. Zhu and X. Yan, "TrustOrch: A Dynamic Trust-Aware Orchestration Framework for Adversarially Robust Multi-Agent Collaboration," 2025.
27. Z. Xu, K. Cao, Y. Zheng, M. Chang, X. Liang and J. Xia, "Generative Distribution Modeling for Credit Card Risk Identification under Noisy and Imbalanced Transactions," 2025.
28. H. Wang, C. Nie and C. Chiang, "Attention-Driven Deep Learning Framework for Intelligent Anomaly Detection in ETL Processes," 2025.
29. J. Lai, C. Chen, J. Li and Q. Gan, "Explainable Intelligent Audit Risk Assessment with Causal Graph Modeling and Causally Constrained Representation Learning," 2025.
30. C. F. Chiang, D. Li, R. Ying, Y. Wang, Q. Gan and J. Li, "Deep Learning-Based Dynamic Graph Framework for Robust Corporate Financial Health Risk Prediction," 2025.
31. U. Pigorsch and S. Schäfer, "High-dimensional stock portfolio trading with deep reinforcement learning," Proceedings of the 2022 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFEr), pp. 1-8, 2022.
32. W. Nuipian, P. Meesad and M. Maliyaem, "Innovative Portfolio Optimization Using Deep Q-Network Reinforcement Learning," Proceedings of the 2024 8th International Conference on Natural Language Processing and Information Retrieval, pp. 292-297, 2024.
33. M. Alidousti, M. Khakzar Bafruei and A. H. Afshar Sedigh, "Pdqn: A Novel Data-Efficient Double Deep Q-Network Framework for Intelligent Financial Portfolio Management," Available at SSRN 5288029, 2025.
34. Y. Huang, K. Cui, Y. Song, et al., "A multi-scaling reinforcement learning trading system based on multi-scaling convolutional neural networks," Mathematics, vol. 11, no. 11, p. 2467, 2023.
35. D. Karzanov, R. Garzón, M. Terekhov, et al., "Regret-Optimized Portfolio Enhancement through Deep Reinforcement Learning and Future Looking Rewards," arXiv preprint arXiv:2502.02619, 2025.
36. S. Yang, "Deep reinforcement learning for portfolio management," Knowledge-Based Systems, vol. 278, p. 110905, 2023.