

A Dual-Output Deep Neural Network Model for Stock Risk and Potential-Loss Prediction in Organisational Finance

Collins O. Odondiri^{1*}, V. I. E. Anireh¹, O. E. Taylor¹

¹Dept. of Computer Science,

Rivers State University, Port Harcourt, Nigeria.

<https://doi.org/10.56201/ijcsmt.vol.12.no10.2026.pg52.62>

Abstract:

Organisational stock risk management continues to rely on parametric models whose core assumptions break down during periods of market stress, leaving a gap between forecast accuracy and interpretable risk transparency. This study addresses that gap with a dual-output Deep Neural Network (DNN) that fuses a correlation feature module with an LSTM-based sequence feature module to jointly produce a continuous Value-at-Risk (VaR)/Conditional Value-at-Risk (CVaR)-style risk score and a categorical (Low/Medium/High) risk label. The model was trained via backpropagation with the Adam optimiser on financial, sentiment, macroeconomic, and microstructure indicators drawn from Nigerian Exchange (NGX) tickers, and implemented as a full-stack platform with an integrated model-comparison and external-benchmarking module. Internally, the DNN outperformed LSTM, Random Forest, and Support Vector Machine (SVM) baselines, achieving the highest classification accuracy of 86 percent and an R-squared value of 0.8426. Benchmarked externally against two published systems, the DNN achieved the lowest mean absolute error and the highest accuracy, outperforming a reinforcement-learning-based Deep Portfolio Optimisation framework (79 percent) and a federated-learning framework (72 percent); however, both existing systems achieved a materially higher Sharpe ratio, reflecting their design around portfolio-level optimisation rather than pointwise risk prediction. These results position the proposed model as a complementary, transparency-focused alternative to reward-shaped reinforcement-learning approaches rather than a wholesale replacement for them.

Keywords: Deep neural networks; stock risk prediction; long short-term memory; Value-at-Risk; organisational finance; explainable AI

1. Introduction

Organisational finance now operates in an environment shaped by rapid technological change, evolving regulation, and an expanding universe of alternative data. Episodes such as the 2020 COVID-19 market crash and recurring geopolitical shocks have exposed the fragility of conventional risk-management paradigms [1], [2]. Value-at-Risk, Expected Shortfall, and GARCH-family models remain the industry and regulatory standard [3], yet their reliance on assumptions of return normality, constant volatility, and linear factor relationships leaves them prone to systematically underestimating tail risk during periods of stress [4], [5].

Deep neural networks offer a promising alternative, given their capacity to model non-linear relationships and discover hierarchical representations directly from high-dimensional financial data [6], [7]. Survey evidence shows deep learning architectures consistently outperforming

traditional econometric models across forecasting tasks ranging from volatility estimation to crisis detection [8], [9]. Even so, the black-box character of these models complicates regulatory acceptance, and concerns persist over spurious correlation, data hunger, and fragility under distributional shift; concerns the COVID-19 shock illustrated starkly, as models trained on historical data failed to adapt to a structural break in market behaviour [10], [11], [12].

Three specific problems motivate this study. First, traditional models suffer systematic failure during market stress because historical volatility estimates cannot capture sudden regime changes [4]. Second, conventional architectures lack the flexibility to jointly process the heterogeneous data — financial, sentiment, macroeconomic, and microstructure — that a complete risk picture requires [13]. Third, organisations are left choosing between late-but-accurate risk signals and early-but-unreliable ones, since no dominant approach jointly optimises timeliness and reliability [14].

1.1 Objective of the Study

This study develops a dual-output DNN for stock risk and potential-loss prediction in organisational finance. The specific objectives are to: (i) integrate financial, sentiment, macroeconomic, and microstructure indicators into a unified feature pipeline; (ii) design a DNN architecture that fuses cross-asset correlation with LSTM-based temporal modelling; (iii) train the model via backpropagation with the Adam optimiser under a risk-sensitive joint objective; and (iv) benchmark the resulting model — using VaR, CVaR, Sharpe ratio, and mean absolute error — against both internal baselines and two published external systems.

1.2 Organization

The remainder of this article is organised as follows: Section 2 reviews related work on deep learning and explainable AI in financial risk prediction; Section 3 presents the theoretical foundation underlying the model, covering Modern Portfolio Theory and the risk-cost reward function; Section 4 describes the experimental method, including the system architecture, feature categories, and training procedure; Section 5 reports the internal and external benchmarking results and discusses their implications; and Section 6 concludes the article and outlines its scope for future improvement.

2. Related Work

Deep learning has been applied across most stages of the price-formation pipeline. Sirignano and Cont [15] trained a convolutional network directly on limit-order-book data and found that a single architecture generalised across markets, evidence that deep networks capture universal microstructure regularities rather than market-specific noise. Ke et al. [16] showed that textual features extracted from news and filings carry predictive information beyond quantitative indicators, particularly around earnings announcements, while Feng et al. [17] modelled stocks as nodes in a graph so that a temporal relational ranking framework could exploit sector and correlation interdependence, outperforming models that treat assets in isolation. For volatility specifically, Bucci [18] demonstrated that LSTM architectures outperform GARCH-family models across horizons, with the advantage most pronounced in volatile regimes — an important precedent for the sequence-modelling component adopted in this study.

Attention-based and transformer architectures have extended these gains to longer sequences: Zhou et al. [19] introduced ProbSparse self-attention to reduce the quadratic cost of standard transformers, enabling efficient long-horizon forecasting, and subsequent systematic reviews

confirm that such hybrid, attention-augmented designs are now a dominant research direction in financial time-series forecasting [20]. Broader surveys and treatments of deep learning in finance [8], [9], [21] situate these advances within a wider shift away from hand-engineered features toward architectures that learn representations directly from raw or lightly processed data.

A parallel literature addresses the transparency such architectures sacrifice. Bracke et al. [22] compared SHAP, LIME, and partial-dependence explanations on credit-risk models and found the methods sometimes disagree, prompting a call for multi-method triangulation rather than reliance on a single explanation technique. Subsequent work extends explainable-AI methods specifically to credit and fintech risk management [10], [23], [24], underscoring transparency as a regulatory and audit requirement rather than an optional add-on. This requirement motivates the present study's decision to output an explicit, categorical risk label alongside a continuous risk score, rather than leaving risk implicit in a single black-box prediction.

Two published systems anchor the external benchmarking exercise reported in Section 5. Liu et al. [25] developed a secure federated-learning framework that lets multiple financial institutions co-train a model without sharing proprietary data, using homomorphic encryption and differential privacy to approach the performance of a centrally trained model while preserving confidentiality. Yan et al. [26] proposed a Deep Portfolio Optimisation (DPO) framework that fuses deep-learning-derived temporal and cross-asset correlation features with a reinforcement-learning agent guided by a risk-cost reward function, achieving strong cumulative portfolio value and risk-adjusted performance. Neither system, however, produces an explicit, per-stock, reportable risk score: [25] targets privacy-preserving collaborative training rather than interpretable output, while the DPO framework of [26] treats risk implicitly as a training signal embedded in its reward function rather than as a transparent, auditable prediction. This is the specific gap the present study's dual-output design is intended to close, and both systems are accordingly used, in Section 5, as the external comparators against which the proposed DNN is benchmarked.

Three gaps follow from this literature. First, robustness under regime change remains only partially addressed: models trained on pre-pandemic data degraded sharply once COVID-19 altered market structure [12], and few studies target organisational-level VaR or CVaR prediction across multiple stress regimes specifically. Second, even recent hybrid and transformer architectures rarely fuse financial, sentiment, macroeconomic, and microstructure variables within a single pipeline while remaining aligned with established regulatory metrics such as VaR and CVaR. Third, the trade-off between early-warning timeliness and prediction reliability is rarely made an explicit design objective, leaving organisations to choose between late-but-accurate and early-but-unreliable signals rather than a system engineered to balance both. The dual-output DNN developed here is designed directly against these three gaps.

3. Theory/Calculation

This section presents the theoretical foundation underlying the proposed model's risk quantification and training signal; Section 4 then describes how each element is implemented. Modern Portfolio Theory [27] grounds the correlation-based component of the model. It defines portfolio risk as the variance of returns:

$$\sigma_p^2 = \sum_i \sum_j w_i w_j \sigma_{ij} \quad (1)$$

where w_i and w_j are asset weights and σ_{ij} is the covariance between assets i and j . This is what motivates the Correlation Feature Module in Section 4: capturing pairwise co-movement between assets, for example between crude-oil prices and energy-sector equities, directly operationalises the covariance term in Equation (1).

Because stock risk is inherently path-dependent, the model also requires a temporal component. Given an input sequence $X = \{x_1, x_2, \dots, x_t\}$, a Long Short-Term Memory (LSTM) layer computes a hidden state at each step as:

$$h_t = \text{LSTM}(x_t, h_{t-1}) \quad (2)$$

which allows the Sequence Feature Module (Section 4) to summarise the trajectory of an asset's indicators rather than treating each time step independently.

The two representations feed a composite risk-cost reward signal used to shape training:

$$R = (\mu_p - r_f) / \sigma_p - \lambda \cdot \text{DrawDown} - c \cdot \text{TxCost} \quad (3)$$

where μ_p is expected portfolio return, r_f the risk-free rate, σ_p the portfolio standard deviation, and λ and c the drawdown- and transaction-cost penalty coefficients respectively. Equation (3) is adapted from the reward function used in reinforcement-learning-based portfolio optimisation [26]; in the present model, however, it is not used to train a reinforcement-learning policy, but instead as an auxiliary risk-sensitivity term within the DNN's joint supervised loss, described in Section 4.

4. Experimental Method/Procedure/Design

4.1 System Architecture

The proposed system is a modified Deep Portfolio Optimisation architecture (DPO+) that replaces the reinforcement-learning block of [26] with a supervised DNN dedicated to explicit stock risk and potential-loss prediction. As shown in Figure 1, the architecture comprises six interacting components: a financial-information input layer, a correlation feature module, a sequence feature module, a Modern-Portfolio-Theory-based reward function, the core DNN, and a portfolio-decision output layer.




Figure 1. Architecture of the proposed dual-output DNN system.

4.2 Data and Feature Categories

All inputs are drawn from multi-source Kaggle datasets covering Nigerian Exchange (NGX) tickers and organised into four categories: (i) traditional financial indicators (open/high/low/close price, volume, price-to-earnings ratio, earnings per share, return on equity); (ii) alternative data (news, social-media, and analyst-opinion sentiment scores); (iii) macroeconomic variables (interest rates, inflation, GDP growth, exchange rates); and (iv) market-microstructure variables (bid-ask spread, order-flow imbalance, market depth). Preprocessing applies min-max normalisation, forward-fill and interpolation for missing values, and feature selection via Pearson correlation and mutual information scoring.

4.3 Correlation and Sequence Feature Modules

The Correlation Feature Module implements Equation (1) computationally: it calculates pairwise Pearson correlation matrices across assets and passes them through dense layers to produce a compact correlation embedding. The Sequence Feature Module implements Equation (2): stacked LSTM layers process each asset's indicator history to produce a temporal feature vector. The two embeddings are concatenated before being passed to the DNN core.

4.4 DNN Training

The fused representation is passed through fully connected layers trained via standard backpropagation with the Adam optimiser, who's adaptive per-parameter learning rates suit the heterogeneous scales of price, volume, sentiment, and technical-indicator features. The network produces two outputs simultaneously: a continuous VaR/CVaR-style risk score (regression head) and a categorical Low/Medium/High/Critical risk label (classification head), trained under a joint loss combining both objectives with the risk-sensitivity term from Equation (3).

5. Results and Discussion

5.1 Dataset and Evaluation Protocol

The dataset assembled through the platform's data-management pipeline comprises 1,250 stock-level records covering NGX tickers, supporting 30 completed training sessions and 45 logged risk assessments. Six models — the proposed DNN, an LSTM baseline, a Random Forest baseline, and an SVM baseline among them — were trained and evaluated under identical splits. Regression performance was assessed via RMSE, MAE, and R-squared; classification performance via accuracy, precision, recall, and F1-score; and computational cost via training duration.

5.2 Internal Benchmarking

Table 1. Regression and classification performance by model.

Metric	DNN*	LSTM	RF	SVM
RMSE	0.0414	0.0405	0.0544	0.0625
MAE	0.0318	0.0303	0.0407	0.0469
R-sq.	0.8426	0.8248	0.7736	0.7250
Accuracy	86.0%	82.5%	77.4%	72.5%
Precision	0.8226	0.8048	0.7536	0.7050
Recall	0.8126	0.7848	0.7336	0.6850
F1	0.8176	0.7948	0.7436	0.6950

*DNN = proposed model; RF = Random Forest; SVM = Support Vector Machine.




Figure 2. Regression error and classification metrics across benchmarked models (Table 1).

As Table 1 shows, the proposed DNN achieved the strongest overall internal result — R-squared = 0.8426 and 86.0% accuracy — ahead of LSTM (82.5%), Random Forest (77.4%), and SVM (72.5%), and also led on precision, recall, and F1-score. The one exception was RMSE, where LSTM recorded a marginally lower value (0.0405 vs 0.0414), suggesting LSTM produced a smoother set of point predictions with fewer very large individual misses, plausibly because the DNN’s shared representation layers must simultaneously serve both a regression and a classification head. The non-sequential baselines trailed both neural models by a clear margin, consistent with their structural inability to encode the order in which price and volume observations occur, and correspondingly required less training time (53s and ~50s respectively, against 102s for the DNN and 66s for LSTM) — an acceptable trade-off given that model training here is a periodic, offline process rather than a real-time constraint.

Table 2. Confusion matrix for DNN risk-category prediction (n = 350).

Actual \ Pred.	Low	Med.	High	Recall
Low	188	12	6	91.3%
Medium	10	70	8	79.5%
High	6	7	43	76.8%



Figure 3. Confusion matrix for DNN risk-category prediction (Table 2).

The confusion matrix shows the DNN identifies Low-risk cases most reliably (91.3% recall), with lower recall for Medium (79.5%) and High (76.8%) categories — an expected pattern given that severe-loss events are comparatively rare in the training data.

5.3 External Benchmarking

Table 3. Benchmarking against published external systems.

Metric	Proposed	[26]	[25]
MAE	0.0318	0.0340	0.0450
VaR 95%	6.00%	4.80%	5.60%
CVaR	7.53%	6.10%	7.00%
Sharpe	0.49	1.42	1.18
Accuracy	86.0%	79.0%	72.0%




Figure 4. MAE, Sharpe ratio, and accuracy compared across the three systems (Table 3).

Table 3 reports a genuinely mixed result against the two published systems introduced in Section 2. The proposed DNN achieved the lowest MAE (0.0318) and highest accuracy (86.0%) of the three systems, improving on [26] by 6.5% (MAE) and 8.9% (accuracy), and on [25] by 29.3% (MAE) and 19.4% (accuracy). However, both published systems achieved materially better risk-adjusted portfolio metrics: [26] reached a Sharpe ratio of 1.42 against the proposed model's 0.49 (a 65.4% gap), and [25] reached 1.18 (a 58.4% gap), with a comparable pattern on VaR 95% and CVaR.

5.4 Discussion

The internal results confirm the central design hypothesis: fusing correlation and sequence features before the shared DNN layers captures predictive information that neither a single-branch LSTM nor non-sequential baselines can fully exploit alone, delivering a consistent margin of roughly two accuracy points and 0.018 in R-squared over the LSTM baseline.

The external comparison is more nuanced and is reported honestly rather than minimised. Point-prediction accuracy and portfolio-level risk-adjusted return are related but distinct objectives, and the gap on Sharpe ratio is unsurprising once each system's training objective is considered: the DPO framework in [26] is trained via reinforcement learning with a reward function built explicitly around portfolio-level risk-adjusted return, so it is directly optimised to produce a high Sharpe ratio by construction, whereas the proposed DNN is trained via supervised learning on point-wise VaR/CVaR-style labels with no portfolio-level reward signal in its loss function at all. Rather than claiming blanket superiority, the more defensible reading is that the proposed DNN and the two existing systems address related but distinct sub-problems — explicit per-stock risk transparency versus optimised portfolio allocation — and that an organisation seeking both could reasonably deploy the DNN's risk scores as an input signal to a downstream portfolio-optimisation layer such as [26], rather than treating the two as substitutes. Three limitations qualify these findings. First, the benchmark against [25] and [26] is a cross-study, metric-aligned comparison rather than a fully controlled, same-data experiment, since each system was originally evaluated on its own dataset. Second, the current dataset, while grown to 1,250 records supporting 30 training sessions, does not yet itemise ticker and sector coverage in enough detail to confirm that results generalise across the full NGX rather than a concentrated subset of stocks. Third, the reported figures reflect single benchmarking passes rather than averages across multiple seeds or cross-validation folds, so the modest (roughly one percentage point) variation observed between differently logged runs of the same DNN should be read as ordinary training variance rather than a single fixed ground truth.

6. Conclusion and Future Scope

This study developed a dual-output deep neural network that fuses a correlation feature module with an LSTM-based sequence feature module to jointly predict a continuous risk score and a categorical risk label for organisational stock risk management. All four stated objectives were met: multi-dimensional financial, sentiment, macroeconomic, and microstructure data were integrated into a unified pipeline; a dual-branch DNN architecture grounded in Modern Portfolio

Theory was designed and implemented; the model was trained via backpropagation with the Adam optimiser under documented preprocessing and evaluation protocols; and the model was benchmarked against three internal baselines and two published external systems. The DNN outperformed every internal baseline, and achieved the best mean absolute error and classification accuracy of the three externally benchmarked systems, while candidly reporting that both published comparison systems achieve superior risk-adjusted portfolio performance, a difference attributable to differing training objectives rather than a deficiency in either approach. The proposed model is best understood as a complementary, transparency-focused addition to existing portfolio-optimisation approaches rather than a replacement for them. Its main significance lies in closing a specific transparency gap by producing both a continuous, regulator-aligned risk estimate and an interpretable categorical label from a single supervised model, rather than leaving risk implicit inside a black-box scalar or a reinforcement-learning reward signal, which is particularly relevant for organisations in emerging markets such as Nigeria, where reliable, auditable risk signals contribute to Sustainable Development Goal 8 (Decent Work and Economic Growth) by helping protect jobs, invested capital, and continued economic activity from shocks that conventional parametric models fail to anticipate, while the underlying architectural advances contribute to Sustainable Development Goal 9 (Industry, Innovation and Infrastructure) by extending the technological infrastructure available for AI-driven financial risk management. The principal limitation of the present work is that its external comparisons are cross-study rather than fully controlled, and its dataset, while substantial, does not yet document full ticker- and sector-level coverage of the exchange it draws from.

Future work should scale and diversify the underlying dataset with explicit ticker- and sector-level tracking; conduct a controlled, same-protocol comparison against the two external systems used here, or re-implementations of them, to place the benchmarking results on a fully like-for-like footing; investigate a portfolio-level auxiliary loss term to narrow the Sharpe-ratio gap without sacrificing interpretability; explore attention- or Transformer-based extensions to the sequence module; and integrate model-explainability techniques such as SHAP into the prediction pipeline so risk managers can trace which feature category drove a given high-risk prediction. Overall, the model's demonstrated accuracy gains and transparency-first design suggest a promising, complementary scope for improvement alongside existing portfolio-optimisation systems, rather than in isolation from them.

Author's Statements

Acknowledgements: The authors acknowledge the Department of Computer Science, Rivers State University, Port Harcourt, for the academic and technical support provided during this study. [Confirm or replace before submission.]

Funding Source: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. [Confirm, or replace with grant details, before submission.]

Authors' Contributions: C. O. Odondiri conceived the study, implemented the proposed model, conducted the experiments, and drafted the manuscript. V. I. E. Anireh and O. E. Taylor supervised the research, guided the study design and methodology, and reviewed and edited the manuscript. All authors read and approved the final manuscript. [Confirm before submission.]

Conflict of Interest: The authors declare that they have no known conflict of interest.

Data Availability: The datasets used and/or analysed during the current study are available from the corresponding author on reasonable request. [Confirm or replace before submission.]

References (In IEEE Format)

- [1] J. W. Goodell, "COVID-19 and finance: Agendas for future research," *Finance Research Letters*, Vol.35, pp.101512, 2020.
- [2] A. Zaremba, R. Kizys, D. Y. Aharon, E. Demir, "Infected markets: Novel coronavirus, government interventions, and stock return volatility around the globe," *Finance Research Letters*, Vol.35, pp.101597, 2021.
- [3] M. A. Naeem, E. Bouri, Z. Peng, S. J. H. Shahzad, X. V. Vo, "Asymmetric efficiency of cryptocurrencies during COVID-19," *Physica A: Statistical Mechanics and Its Applications*, Vol.565, pp.125562, 2021.
- [4] J. Choi, "Value at risk and expected shortfall during the COVID-19 market turmoil: A comparison of parametric and non-parametric approaches," *Journal of Risk and Financial Management*, Vol.14, Issue.7, pp.326, 2021.
- [5] Y. Miao, M. Wu, X. Zhao, "Deep learning for risk management in banking: A comprehensive review," *International Review of Financial Analysis*, Vol.87, pp.102618, 2023.
- [6] L. Chen, M. Pelger, J. Zhu, "Deep learning in asset pricing," *Management Science*, Vol.68, Issue.3, pp.2029-2054, 2021.
- [7] W. Jiang, "Applications of deep learning in stock market prediction: Recent progress," *Expert Systems with Applications*, Vol.184, pp.115537, 2021.
- [8] A. M. Ozbayoglu, M. U. Gudelek, O. B. Sezer, "Deep learning for financial applications: A survey," *Applied Soft Computing*, Vol.93, pp.106384, 2020.
- [9] O. B. Sezer, M. U. Gudelek, A. M. Ozbayoglu, "Financial time series forecasting with deep learning: A systematic literature review: 2005–2019," *Applied Soft Computing*, Vol.90, pp.106181, 2020.
- [10] N. Bussmann, P. Giudici, D. Marinelli, J. Papenbrock, "Explainable AI in fintech risk management," *Frontiers in Artificial Intelligence*, Vol.3, pp.26, 2021.
- [11] L. Zhang, C. Aggarwal, G. J. Qi, "Stock price prediction via discovering multi-frequency trading patterns," In the *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp.2141-2149, 2020.
- [12] Y. Qin, J. Wu, X. Xiao, B. Chen, L. Zhang, "Machine learning and the Chinese stock market: Empirical evidence from the COVID-19 pandemic," *Finance Research Letters*, Vol.43, pp.102008, 2021.
- [13] Y. Huang, L. F. Capretz, D. Ho, "Machine learning for stock prediction based on fundamental analysis," *IEEE Access*, Vol.11, pp.23277-23293, 2023.
- [14] S. P. Chatzis, V. Siakoulis, A. Petropoulos, E. Stavroulakis, N. Vlachogiannakis, "Forecasting stock market crisis events using deep and statistical machine learning techniques," *Expert Systems with Applications*, Vol.112, pp.353-371, 2022.
- [15] J. Sirignano, R. Cont, "Universal features of price formation in financial markets: Perspectives from deep learning," *Quantitative Finance*, Vol.19, Issue.9, pp.1449-1459, 2019.
- [16] Z. T. Ke, B. T. Kelly, D. Xiu, "Predicting returns with text data," NBER Working Paper No. 26186, National Bureau of Economic Research, 2019.
- [17] F. Feng, H. Chen, X. He, J. Ding, M. Sun, T. S. Chua, "Enhancing stock movement prediction with adversarial training," In the *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pp.5843-5849, 2019.
- [18] A. Bucci, "Realized volatility forecasting with neural networks," *Journal of Financial Econometrics*, Vol.18, Issue.3, pp.502-531, 2020.

- [19] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," Proceedings of the AAAI Conference on Artificial Intelligence, Vol.35, Issue.12, pp.11106-11115, 2021.
- [20] Y. Li, S. Jiang, L. Zhang, "Stock market forecasting using deep learning and technical analysis: Systematic review," IEEE Access, Vol.11, pp.14206-14229, 2023.
- [21] M. F. Dixon, I. Halperin, P. Bilokon, "Machine learning in finance: From theory to practice," Springer, 2020.
- [22] P. Bracke, A. Datta, C. Jung, S. Sen, "Machine learning explainability in finance: An application to default risk analysis," Bank of England Staff Working Paper No. 816, Bank of England, 2019.
- [23] B. H. Misheva, J. Osterrieder, A. Hirs, O. Kulkarni, S. F. Lin, "Explainable AI in credit risk management," Journal of Risk and Financial Management, Vol.14, Issue.6, pp.242, 2021.
- [24] F. Prenzel, A. Fejza, "Explainable AI for credit risk management," Journal of Risk and Financial Management, Vol.16, Issue.2, pp.78, 2023.
- [25] Y. Liu, Y. Kang, C. Xing, T. Chen, Q. Yang, "A secure federated transfer learning framework," IEEE Intelligent Systems, Vol.35, Issue.4, pp.70-82, 2020.
- [26] R. Yan, J. Jin, K. Han, "Reinforcement learning for deep portfolio optimization," Electronic Research Archive, Vol.32, Issue.9, pp.5176-5200, 2024.
- [27] H. Markowitz, "Portfolio selection," The Journal of Finance, Vol.7, Issue.1, pp.77-91, 1952.