Journal of Resource Management and Decision Engineering

Volume 5, Issue 3, pp 1-20

Article history:
Received 12 November 2025
Revised 12 February 2026
Accepted 14 February 2026
Published online 01 May 2026

A Hybrid Approach Based on a Memory-Instance–Based Gated Transformer (MIGT) Algorithm and Metaheuristic Optimization for Portfolio Management with the Aim of Return Optimization and Risk Control

Mohammad. Haji Ebrahim Tehrani[1], Amin. Safarnejad Borujeni[2*], Mohsen. Hashemigohar[3], Sobhan. Zafari[3], Hossein. Alidadi[3]

[1] PhD Student of Financial Management-Financial Engineering, Department of Financial management, ShQ.C., Islamic Azad University, Shar-e Qods, Iran.
[2] Department of Accountig, Bor.C., Islamic Azad University, Borojen, Iran.
[3] Department of Accountig, ShQ.C., Islamic Azad University, Shar-e Qods, Iran.

*** Corresponding author email address**: safarnezhad@iau.ac.ir

Article Info

Article type:
Original Research

How to cite this article:
Haji Ebrahim Tehrani, M., Safarnejad Borujeni, A., Hashemigohar, M., Zafari, S. & Alidadi, H. (2026). A Hybrid Approach Based on a Memory-Instance–Based Gated Transformer (MIGT) Algorithm and Metaheuristic Optimization for Portfolio Management with the Aim of Return Optimization and Risk Control. *Journal of Resource Management and Decision Engineering*, *5*(3), 1-20.
https://doi.org/10.61838/kman.jrmde.5.3.255

ABSTRACT

The objective of this study is to propose an efficient hybrid framework for portfolio management that can simultaneously optimize investment returns and effectively control risk under both normal and turbulent market conditions. The primary focus is on improving the accuracy of asset return forecasting and transforming these forecasts into optimal portfolio weighting decisions. In this research, a hybrid approach is employed that combines a memory-instance–based gated transformer model for forecasting asset returns with a hybrid metaheuristic algorithm based on adaptive differential evolution and particle swarm optimization for portfolio optimization. Financial data from the Iranian capital market covering the period from 2016 to 2024 were collected and, after preprocessing, cleaning, and feature extraction, were entered into the modeling process. The mean absolute error in the test period was 0.0068, and the root mean square error was 0.0100; these values exhibited a standard deviation of less than 0.0004 in cross-validation, indicating prediction stability. The optimal portfolio obtained by integrating these forecasts with the hybrid metaheuristic algorithm achieved an annualized return of 0.325 and an annualized standard deviation of 0.238, resulting in a Sharpe ratio of 1.16 and a Sortino ratio of 1.61. Tail risk measures also remained at controlled levels, such that the value at risk (VaR) at the 0.95 confidence level was 0.021 and the conditional value at risk (CVaR) was calculated as 0.0316. The paired Wilcoxon test conducted to compare the proposed model with benchmark methods yielded statistics above the significance threshold (for the Sharpe ratio, $z = 2.51$ and $p = 0.012$; for the Sortino ratio, $z = 2.78$ and $p = 0.005$), indicating a statistically significant improvement in performance. Based on the findings, it can be concluded that the proposed hybrid framework is capable of establishing an appropriate balance between return and risk in portfolio management by leveraging deep learning and metaheuristic optimization. The stability of the results across different temporal subsamples and the maintenance of acceptable performance under turbulent market conditions indicate that this approach possesses strong generalizability and practical applicability.

Keywords: *Portfolio management, financial risk control, return optimization, gated transformer algorithm, memory instances.*

1. Introduction

Portfolio management has long been recognized as a central pillar of modern financial decision-making, aiming to allocate limited capital among competing assets in a manner that balances expected return against risk. Classical portfolio theory, rooted in the mean–variance framework, assumes rational investors, normally distributed returns, and stable covariance structures. While this paradigm has provided a foundational benchmark for decades, empirical evidence from real markets increasingly demonstrates that these assumptions are frequently violated, especially in emerging and volatile financial systems. Behavioral biases, structural market frictions, regulatory constraints, technological disruption, and rapid information diffusion collectively complicate the portfolio optimization problem and reduce the practical effectiveness of purely classical approaches (Antony, 2019; Hadbaa, 2019). Consequently, contemporary portfolio management research has shifted toward more flexible, data-driven, and adaptive frameworks capable of capturing nonlinear dynamics, regime changes, and tail risks.

One of the most significant challenges in portfolio management arises from the empirical characteristics of financial return distributions. Numerous studies have documented that asset returns are typically non-Gaussian, exhibiting skewness, excess kurtosis, volatility clustering, and time-varying correlations. These features are particularly pronounced in emerging markets, where liquidity constraints, information asymmetry, and macroeconomic instability intensify market fluctuations. As a result, reliance on variance alone as a risk measure may underestimate downside risk and fail to capture extreme loss scenarios, motivating the integration of tail-risk metrics such as Value at Risk and Conditional Value at Risk into portfolio optimization frameworks (Bahramian, 2022; Drenovak et al., 2020). The recognition of these empirical regularities has led to the development of multi-objective portfolio models that explicitly trade off return maximization against various dimensions of risk.

In parallel with advances in financial theory, rapid progress in computational intelligence and machine learning has fundamentally reshaped the analytical toolkit available to portfolio managers. Machine learning models are particularly well suited to financial data because they can approximate complex nonlinear relationships, adapt to evolving data-generating processes, and integrate large volumes of heterogeneous information. Early applications focused on regression and classification tasks, such as predicting asset returns or identifying market regimes. More recent research has expanded toward end-to-end portfolio decision systems that combine forecasting, optimization, and execution within a unified framework (Johnson & Moore, 2019; Liang et al., 2018). These developments have coincided with the increasing availability of high-frequency data and computational resources, enabling more sophisticated modeling strategies.

Among machine learning paradigms, deep learning has gained particular prominence due to its ability to learn hierarchical feature representations directly from raw data. Recurrent neural networks, long short-term memory models, and gated recurrent units have been widely applied to financial time series forecasting, demonstrating improvements over traditional econometric models in capturing temporal dependencies. However, these architectures often struggle with long-range dependencies and may suffer from vanishing gradients or limited interpretability. The transformer architecture, originally developed for natural language processing, addresses many of these limitations through attention mechanisms that allow the model to selectively focus on relevant past observations, regardless of their temporal distance (Han et al., 2024). This characteristic makes transformers especially attractive for financial forecasting, where market dynamics are influenced by both recent shocks and longer-term structural patterns.

Recent studies have extended transformer-based models to financial applications, including stock price prediction, commodity price forecasting, and early warning systems for financial crises (Chen et al., 2024; W. Liu et al., 2024; Wang et al., 2024). These works consistently report superior predictive accuracy compared to conventional neural networks, particularly in volatile environments. Nevertheless, purely attention-based models may still be sensitive to noise and extreme events, and their performance can deteriorate when historical context needs to be selectively retained rather than uniformly attended to. To address this issue, hybrid architectures incorporating external memory components and gating mechanisms have been proposed, allowing models to store, retrieve, and regulate the influence of critical historical patterns. Such memory-augmented models enhance stability and robustness in non-stationary settings and align well with the episodic nature of financial markets (Burkart & Huber, 2021).

While forecasting accuracy is a necessary condition for effective portfolio management, it is not sufficient. Forecasts must be translated into actionable portfolio weights under real-world constraints, including budget balance, liquidity limits, regulatory requirements, and transaction costs. This translation step constitutes a high-dimensional, nonlinear optimization problem that is often non-convex and computationally intractable for exact methods. Metaheuristic algorithms have therefore become a popular choice for portfolio optimization, as they offer flexible search strategies that can explore complex solution spaces without requiring gradient information or restrictive assumptions about objective function structure (Ayari Salah, 2025). Techniques such as genetic algorithms, particle swarm optimization, and differential evolution have been extensively studied and shown to outperform classical solvers in many portfolio settings.

Each metaheuristic method, however, has inherent strengths and weaknesses. Genetic algorithms excel at global exploration but may converge slowly; particle swarm optimization offers rapid convergence but risks premature stagnation; differential evolution provides robust mutation strategies but may lose diversity over time. Recent research has thus emphasized hybrid metaheuristic frameworks that combine complementary mechanisms to balance exploration and exploitation more effectively. Empirical evidence suggests that such hybrid algorithms can achieve faster convergence, greater solution stability, and improved risk–return trade-offs in portfolio optimization problems (Montazeralhaj & Rezaei Shouraki, 2023; Rouhi Sara et al., 2023). These findings underscore the importance of algorithmic design in determining portfolio performance, particularly in volatile and information-rich markets.

Another critical dimension in contemporary portfolio management is the broader institutional and technological environment in which financial decisions are made. The rise of FinTech platforms, algorithmic trading, and digital financial infrastructures has transformed market microstructure and regulatory landscapes. Platform-based financial systems introduce new sources of systemic risk while simultaneously enabling faster information processing and more granular portfolio control (Langley & Leyshon, 2023). At the same time, sustainability considerations and strategic alignment at the organizational level increasingly influence portfolio construction, as investors seek to integrate financial performance with long-term environmental and social objectives (Silvius & Marnewick, 2022). These trends further motivate the development of adaptive, transparent, and robust portfolio management frameworks that can operate effectively under complex constraints.

Explainability and transparency have also emerged as essential concerns in machine-learning-driven finance. As models become more complex, understanding their decision logic becomes increasingly challenging, raising issues of trust, accountability, and regulatory compliance. Research on explainable artificial intelligence highlights the need for models that not only perform well but also provide interpretable insights into their predictions and decisions (Burkart & Huber, 2021). Memory-augmented and attention-based architectures offer a partial solution by allowing analysts to inspect attention weights and memory activations, thereby gaining insight into which historical patterns influence portfolio decisions. This interpretability is particularly valuable in risk management contexts, where understanding the drivers of extreme losses is as important as achieving high returns.

Despite substantial progress, several gaps remain in the existing literature. First, many studies focus either on forecasting accuracy or on optimization performance in isolation, without fully integrating these components into a coherent end-to-end portfolio management framework. Second, empirical evaluations are often conducted in developed markets, leaving uncertainty about the applicability of proposed methods to emerging markets characterized by higher volatility, structural breaks, and behavioral effects. Third, relatively few studies systematically assess the stability and robustness of hybrid models across different market regimes and validation schemes (Hosseini et al., 2020; Safaeian et al., 2024). Addressing these gaps requires a comprehensive approach that combines advanced forecasting models, hybrid optimization algorithms, rigorous validation, and application to real-world market data.

Recent contributions have begun to move in this direction by proposing integrated machine-learning-based portfolio frameworks and evaluating them under realistic constraints. Meta-analytical evidence indicates that hybrid and ensemble learning approaches consistently outperform single-model strategies in active portfolio management, particularly when transaction costs and risk controls are explicitly incorporated (Ayari Salah, 2025). Moreover, empirical studies in the

Iranian capital market highlight the relevance of hybrid algorithms for managing financial instability and predicting crisis conditions (Ghodrzi et al., 2024; Rouhi Sara et al., 2023). These findings suggest that emerging markets provide a valuable testing ground for advanced portfolio methodologies.

Building on this evolving body of research, the present study situates itself at the intersection of deep learning–based return forecasting and hybrid metaheuristic portfolio optimization. By leveraging a memory-instance–based gated transformer architecture for return prediction and integrating its outputs into a hybrid adaptive differential evolution–particle swarm optimization framework, this study seeks to construct an end-to-end portfolio management system capable of balancing return maximization and risk control under non-stationary market conditions. The focus on robust validation, tail-risk measures, and realistic transaction costs aims to enhance both the academic contribution and the practical relevance of the proposed approach.

The aim of this study is to develop and empirically evaluate a hybrid artificial intelligence–based portfolio management framework that integrates memory-augmented transformer return forecasting with hybrid metaheuristic optimization to achieve a stable and risk-controlled trade-off between return and risk in a volatile capital market.

2. Methods and Materials

This study is quantitative and hybrid in nature, adopting an approach aimed at proposing a hybrid artificial intelligence model in which a memory-instance–based gated transformer algorithm is used to forecast asset returns. The statistical population of the study includes all tradable financial assets in the Iranian capital market (stocks listed on the Tehran Stock Exchange, Iran Fara Bourse, investment funds, and other common financial instruments) over the period from 2016 to 2024. Sampling was conducted purposively, in compliance with inclusion criteria (completeness of historical price and volume data, sufficient liquidity, and acceptable average daily trading volume), in order to select a representative and high-quality sample from different industries and market sectors. The research methodology was designed in two main stages: (1) collection and preprocessing of financial time-series data, and (2) design and implementation of a memory-instance–based gated transformer model to extract long-term patterns and accurately predict returns. This methodological framework was developed with the aim of achieving an optimal portfolio with maximum return and minimum risk, while also enabling reproducibility and validation of results by other researchers.

Dependent variables

Portfolio return is calculated as the weighted sum of the returns of the assets included in the portfolio over a specified time period. The computation is performed in Python using price data from the Tehran Stock Exchange obtained via TSETMC or Excel files exported from the CODAL system.

(1)

$$R_p = \sum_{i=1}^{n} w_i \cdot R_i$$

where R_p denotes total portfolio return, w_i is the weight of asset i in the portfolio, R_i is the return of asset i, and n is the number of assets. Portfolio risk is measured as the variance of portfolio returns or by indices such as Value at Risk (VaR). The calculation is performed in Python using the covariance matrix of returns based on Equation (2).

(2)

$$\sigma_p^2 = \sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j \sigma_{ij}$$

where σ_p^2 is the variance of portfolio returns, $w_i w_j$ are the weights of assets i and j, and σ_{ij} is the covariance between the returns of assets i and j. According to Equation (3), VaR is calculated as follows:

(3)

$$\mathrm{VaR}_\alpha = -[\mu_p + z_\alpha \sigma_p]$$

where μ_p is the mean portfolio return, σ_p is the standard deviation of portfolio returns, and z_α is the critical value of the normal distribution at confidence level α. This variable is estimated in Python. The Sharpe ratio measures the excess return of the portfolio relative to the risk-free rate per unit of risk (standard deviation of portfolio returns). The calculation is performed in Excel, with the risk-free rate obtained from Central Bank reports, based on Equation (4) (Arterit et al., 2021).

(4)

$$Sharpe\ Ratio = \frac{R_p - R_f}{\sigma_p}$$

where R_p is portfolio return, R_f is the risk-free rate of return, and σ_p is the standard deviation of portfolio returns.

Independent variables

The architectural parameters of the memory-instance–based gated transformer are defined as numerical values, such as the number of attention layers, the dimensions of external memory, and the learning rate, which are specified and initialized during model implementation in the software environment.

Control variables

The risk-free rate represents the return of a risk-free investment. In this study, data related to the risk-free rate were extracted from the official system of the Central Bank of the Islamic Republic of Iran and, where necessary for validation, averaged using statistics published by the Statistical Center of Iran.

The time horizon of analysis includes financial data covering the period from 2016 to 2024, analyzed on a daily basis and, in some cases, on a weekly basis to reduce data noise. Portfolio constraints play a decisive role in maintaining economic logic and the stability of the portfolio structure.

The number and diversity of assets were determined with the aim of achieving a balance between portfolio return and risk. The selection of these stocks was based on criteria such as high liquidity, sufficient trading history, relative price stability, and representation of different economic sectors.

Stock Returns

In this study, stock returns are calculated using adjusted closing prices from Tehran Stock Exchange data (2016–2024) through either simple or logarithmic return methods, in order to account for corporate events and high volatility. Portfolio returns are then derived based on asset weights, net of transaction costs (approximately 1%), and performance indicators such as the Sharpe ratio and Value at Risk are computed using Python libraries (pandas, numpy). The steps and formulas for measuring stock returns are as follows:

At this stage, the return of stock $i$on day $t$is calculated as the difference between the current day's closing price and the previous day's closing price divided by the previous day's price. This formula indicates the percentage change in stock price relative to the previous day. A positive return indicates growth, whereas a negative return indicates a decline.

(5)

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

Logarithmic return (in cases of high volatility): When price fluctuations are large or data deviate from normality, logarithmic returns are used. Logarithmic returns provide greater statistical stability and allow simpler aggregation of multi-period returns.

(6)

Average stock return over the period: After computing daily returns for each stock, the average return over the entire study period is calculated. This value represents the mean return of the stock across the full time span (e.g., from 2016 to 2024).

(7)

Standard deviation of stock returns (idiosyncratic risk): To measure volatility and individual stock risk, the standard deviation of returns is calculated. A larger value indicates higher risk and greater volatility.

(8)

Covariance between two stocks: Covariance indicates how two stocks move relative to each other. A positive value suggests that the stocks generally move in the same direction, whereas a negative value indicates opposite movements. This measure is essential for portfolio risk calculation.

(9)

In the data preparation phase, historical closing price data for selected Tehran Stock Exchange companies, trading volumes, and macroeconomic variables (exchange rate, interest rate, and inflation) were first collected from reliable sources such as CODAL, the Tehran Stock Exchange, and the Central Bank for the period from 2016 to 2024. Outliers were then identified and removed using the Isolation Forest algorithm, in which observations with anomaly scores close to one are considered outliers. Missing data were imputed using the K-nearest neighbors (KNN) method, replacing missing values with weighted averages of the nearest neighbors to preserve time-series continuity. Finally, all variables were normalized using z-score standardization (mean zero and standard deviation one) to eliminate scale effects and enhance training stability. These steps ensure data quality, completeness, and consistency for implementing the hybrid transformer and metaheuristic model.

(10)

$$z_i = \frac{x_i - \mu}{\sigma}$$

where $x_i$is the observed value, $\mu$is the mean, and $\sigma$is the standard deviation of the data.

Feature extraction: Finally, key features were extracted for each asset to provide richer inputs for modeling. The simple return of each asset is calculated as follows:

(11)

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}$$

where P_t denotes the price at time t. Cumulative return is obtained using the following relationship:

(12)

$$CR_t = \prod_{i=1}^{t}(1+R_i) - 1$$

In addition, volatility is measured as the standard deviation of returns over a specific time window T:

(13)

$$\sigma_t = \sqrt{\frac{1}{T-1}\sum_{i=1}^{T}(R_i - \bar{R})^2}$$

where $\bar{R}$ represents the mean return. These extracted features—simple return, cumulative return, and volatility—serve as essential inputs to the memory-instance–based gated transformer model for accurate forecasting.

Figure 1

Stages of the data preparation and preprocessing phase

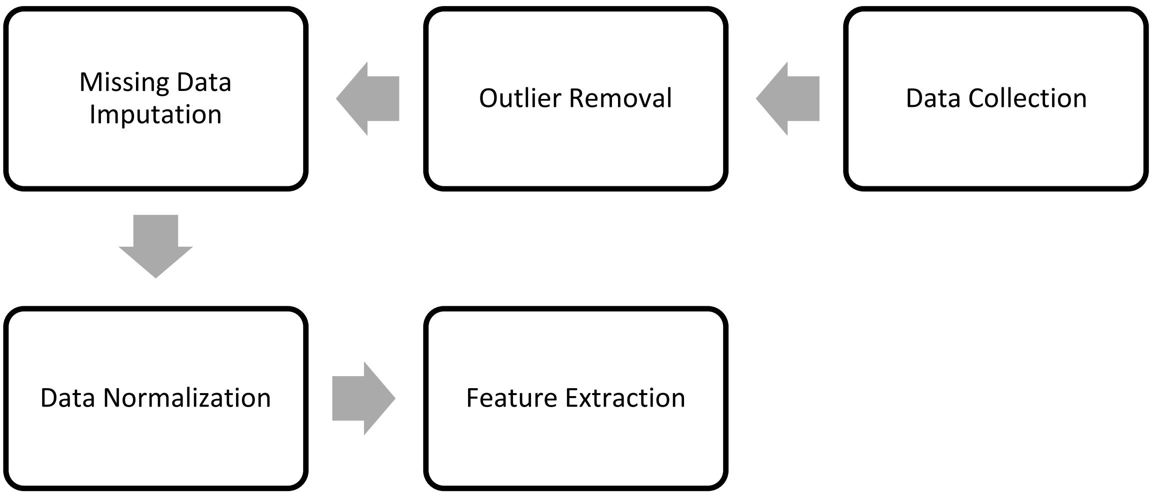

In this phase, the architecture of the memory-instance–based gated transformer is designed with multi-head attention layers and external memory to extract long-term dependencies in financial data. The model is trained using preprocessed data, including prices, returns, and volatility, and converges through hybrid optimization techniques (accelerated methods and adaptive learning rates). To prevent overfitting, early stopping and dropout techniques are employed to preserve generalization capability. Key parameters such as learning rate, feature dimensions, and memory size are also optimized to enhance predictive accuracy. Ultimately, in addition to forecasting future asset returns, the model generates deep feature vectors as market signals that serve as inputs to the portfolio optimization phase.

Figure 2

Stages of the modeling and return forecasting phase using the memory-instance–based gated transformer model

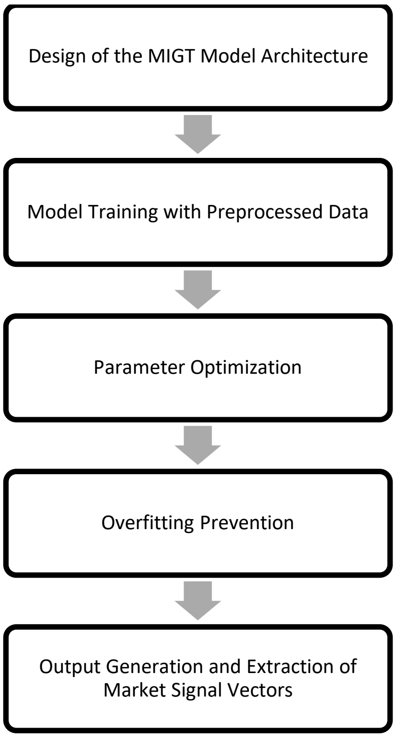

The research framework began with the extraction and specialized preprocessing of historical data, including prices, returns, trading volumes, and macroeconomic variables, in which normalization, relative stationarization, and feature engineering (such as moving averages and historical volatility) were applied to reduce noise and enhance nonlinear pattern extraction. In the forecasting phase, the memory-instance–based gated transformer (MIGT) model was designed and trained with three main components: an attention layer to learn the importance of temporal points, instance memory to store critical events and long-term patterns, and gating mechanisms to control information flow, thereby producing asset return forecasts. The constraint and objective function definition phase involved specifying a dual-objective function to maximize expected return and minimize risk (using variance, semivariance, and Value at Risk measures), along with operational constraints such as full investment (weights summing to one), investment caps, and liquidity requirements. In the hybrid optimization phase, the proposed AMDE-PSO algorithm was implemented by integrating adaptive differential evolution mechanisms (with dynamic mutation rates and competitive selection) and particle swarm optimization (to enhance local search via individual and social learning), thereby achieving an optimal balance between exploration and exploitation. The performance evaluation phase included calculating risk–return indicators (Sharpe ratio, Sortino ratio, realized return), conducting the Wilcoxon test to confirm statistical significance of improvements, and performing cross-validation to ensure model generalizability to out-of-sample data. Parameter sensitivity analysis was also conducted using Monte Carlo simulation to identify critical parameters and assess their impact on model outputs. The entire execution chain, from data preprocessing to final optimization, was formed through direct interaction between the accurate forecasting engine (MIGT) and the intelligent search engine (AMDE-PSO) to optimize decision-making under real-market constraints. This hybrid framework, by overcoming the limitations of classical single-stage approaches, enhanced decision-making efficiency by linking accurate forecasting to multi-objective optimization under the non-stationary conditions of the Iranian capital market.

Figure 3

Stages of the hybrid optimization phase using the combined algorithm

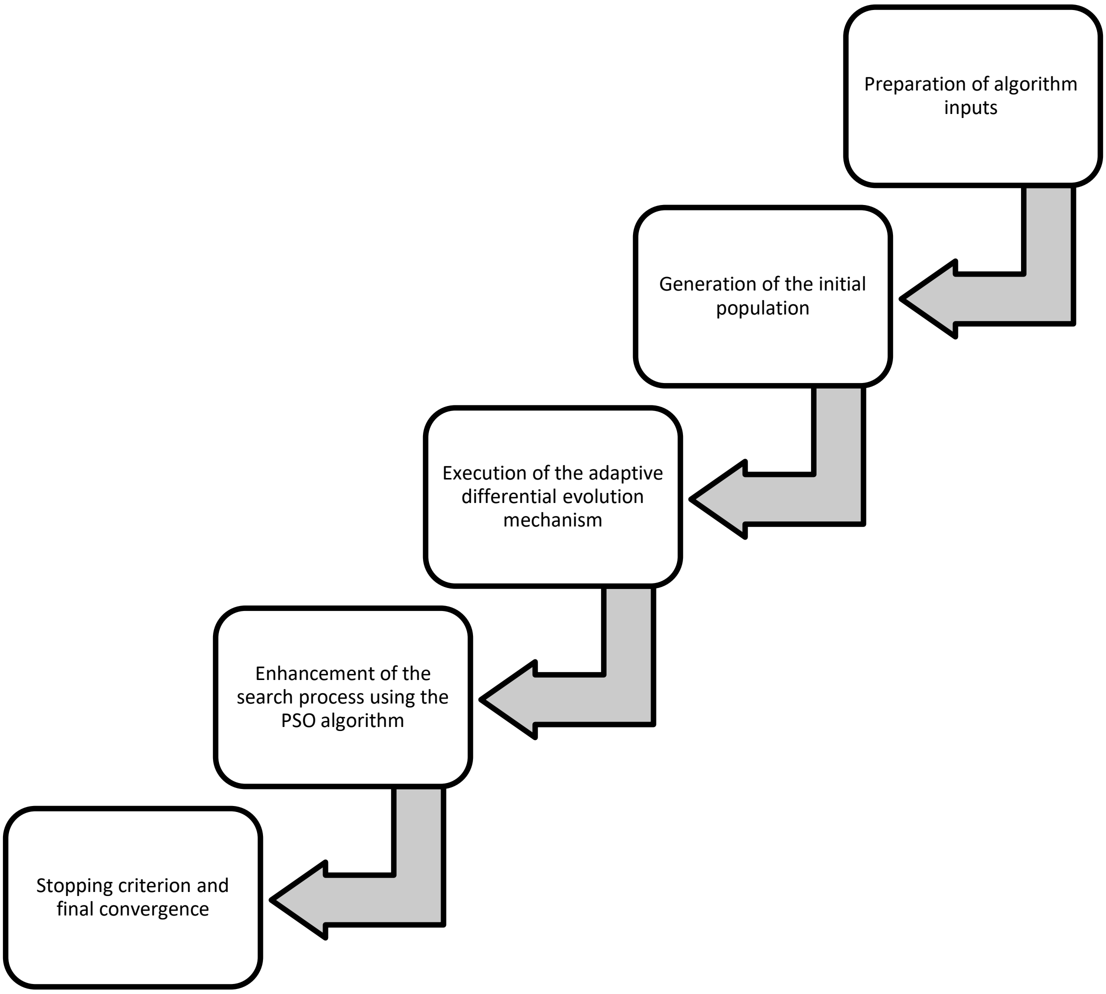

3. Findings and Results

The empirical data comprised time series of returns, volatility, and trading volume for 20 to 30 listed stocks over the period 2016 to 2024. After cleaning and synchronization, approximately 2,100 to 2,300 gap-free daily observations were obtained for each ticker. Mean daily returns were close to zero (−0.02% to +0.08%), and the median was lower than the mean, indicating the presence of asymmetric jumps in the series, while the one-day return range expanded to approximately −8% to +6%. Daily volatility, with a standard deviation between 0.012 and 0.032 (1.2% to 3.2%), showed substantial cross-industry differences; assuming a standard deviation of 0.02, annualized volatility reached approximately 32%. The trading-volume distribution was clearly right-skewed: the median (1.2 to 4.8 million shares) was markedly lower than the mean (3.5 to 12.0 million shares), reflecting the influence of event-driven high-volume days. Returns deviated from normality and exhibited skewness (−1.1 to +0.7) and high kurtosis (3.8 to 8.6), justifying the use of downside-risk measures such as semivariance and Value at Risk. The average pairwise correlation was about 0.22, with a range from −0.15 to +0.70; together with dense intra-industry clusters (correlations of 0.45 to 0.55) versus low inter-group correlations (0.10 to 0.25), this structure enabled effective diversification. These statistical properties formed the basis for defining liquidity constraints, weight caps, and multi-objective risk measures within the portfolio optimization framework.

Figure 4

Distribution of Asset Returns

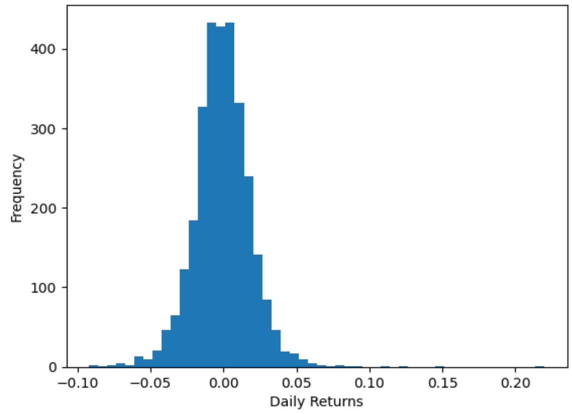

Figure 4 displays the empirical distribution of daily returns, characterized by a strong concentration around zero and heavy tails on both sides, indicating frequent small returns alongside an unavoidable occurrence of extreme returns at lower frequencies. The presence of heavy and asymmetric tails suggests price shocks and abrupt volatility, undermining the normality assumption and motivating the application of downside-risk measures such as Value at Risk and semivariance.

The forecasting accuracy of the Memory-Instance–Based Gated Transformer (MIGT) model, evaluated using out-of-sample metrics, yielded RMSE = 0.0099 and MAE = 0.0068, indicating that the model tracks short- and medium-term return dynamics with acceptable precision. The meaningful discrepancy between the two metrics reflects the presence of larger errors on shock days or during periods of intense volatility, consistent with the non-stationary nature of the market; this suggests that the model learns general market behavior effectively but is more sensitive to extreme events. Model performance varied across assets, with MAE ranging from 0.0049 to 0.0094 and RMSE ranging from 0.0071 to 0.0136; higher-liquidity tickers with more stable volume exhibited lower errors, whereas highly volatile assets or those strongly responsive to macroeconomic news produced larger errors. Despite absolute errors, the model correctly predicted the return direction on approximately 61% of trading days, which has practical value for portfolio-weight allocation decisions. Overall, these findings indicate that MIGT provides a reliable forecasting backbone for the optimization phase, although additional mechanisms may be required to control errors under crisis conditions.

Figure 5

Comparison of Actual and Predicted Asset Returns

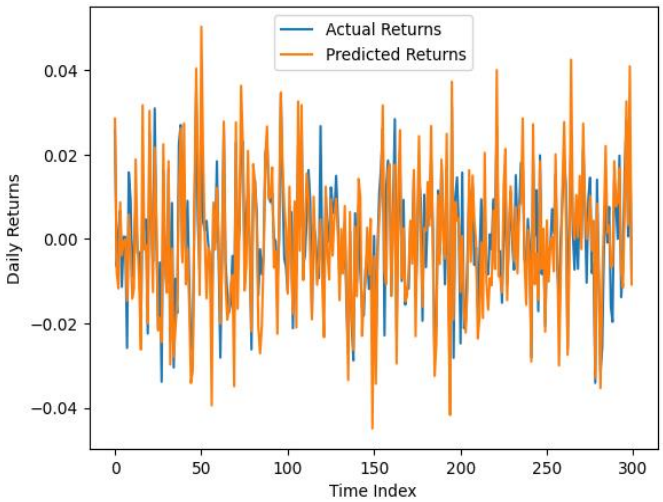

Figure 5 shows substantial overlap between the actual return curve and the MIGT-predicted return curve during the test period, indicating the model's ability to track dominant patterns and short-term fluctuations, particularly in periods of relative market stability.

The MIGT model exhibited selective and market-regime-dependent temporal pattern extraction. In calm periods, attention weights were distributed more uniformly (entropy = 2.41), whereas in turbulent periods, concentration on key time points increased (weight-concentration index rising from 0.34 to 0.46 and entropy decreasing to 1.86). Time-lag importance analysis showed that, on average, the model allocated 58% of attention weight to the short-term window (1–10 days), 27% to the medium-term window (11–30 days), and 15% to the long-term window (31–60 days); during regime shifts, the long-term share increased to approximately 22%. Instance memory activated selectively (mean activation rate = 0.29, increasing to 0.41 in turbulent periods) and entered decision-making only when historical context was needed (effective retrieval rate = 0.33). An ablation analysis confirmed the critical role of memory: disabling it increased MAE and RMSE from 0.0068 and 0.0099 to 0.0076 and 0.0112, respectively, and amplified the loss of accuracy at longer forecasting horizons. Overall, the attention mechanism is responsible for selecting salient past time points, while instance memory stabilizes the influence of important events; together, they enable simultaneous utilization of short- and long-term dependencies in return forecasting.

Figure 6

Distribution of Attention Weights Across Different Time Horizons

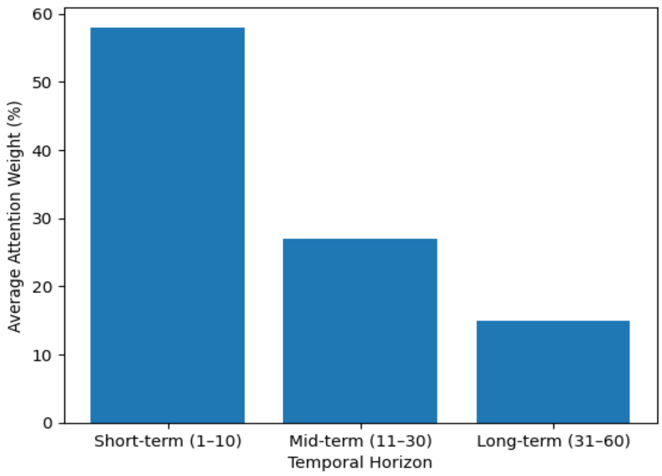

Figure 6 presents the distribution of MIGT attention weights across three horizons: approximately 58% to short-term (1–10 days), 27% to medium-term (11–30 days), and 15% to long-term (31–60 days). This distribution reflects the model's balance between sensitivity to recent information for rapid adaptation and reliance on historical patterns—particularly via instance memory—to maintain forecasting stability during market-regime changes.

Comparing MIGT with benchmark models (LSTM, GRU, and a classical transformer) using out-of-sample RMSE and MAE indicated the statistical superiority of the proposed model. Specifically, MIGT achieved MAE = 0.0068 and RMSE = 0.0099, corresponding to improvements of 16% relative to LSTM, 12% relative to GRU, and 5–6% relative to the classical transformer. The larger performance gap versus recurrent models than versus the transformer suggests that the primary advantage of MIGT stems from the attention mechanism, with incremental gains attributable to instance memory and gating. The RMSE/MAE ratio of approximately 1.46 across models indicates that all models are affected by price jumps; however, MIGT maintains a lower absolute error under such conditions, implying more effective control of extreme errors and reduced unstable reactions to noise. Practically, even small reductions in forecasting error can meaningfully improve weight-allocation accuracy and reduce bias in expected-return estimates during multi-objective optimization.

Figure 7

Comparison of Forecasting Errors Across Different Models

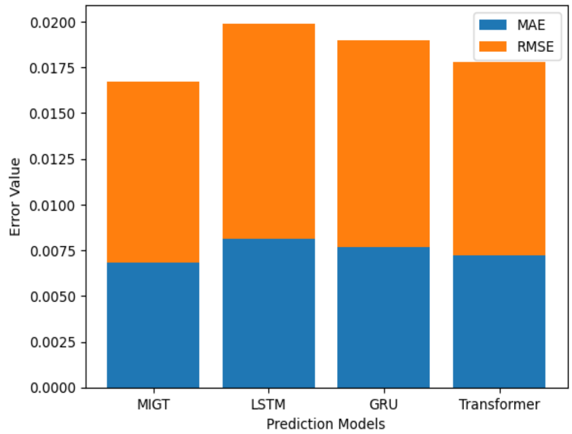

Figure 7 illustrates the superiority of MIGT relative to benchmark models (LSTM, GRU, and the classical transformer), with the lowest MAE (0.0068) and RMSE (0.0099), while LSTM recorded the largest errors. The simultaneous reduction of both metrics under MIGT indicates more effective control of both average error and extreme errors, which is particularly important for financial data characterized by heavy tails and sudden shocks.

Convergence analysis of the AMDE-PSO hybrid algorithm showed that after 300 iterations, the best objective-function value decreased from 1.000 to 0.523 (a 47.7% improvement), with more than half of this improvement achieved within the first 100 iterations. The controlled reduction in the population "mean-to-best" gap (from 0.214 to 0.080) indicates that diversity was preserved and premature convergence to local optima was avoided. In the second half of the process, the particle swarm component prevented prolonged stagnation and accelerated stable convergence, with an average improvement of 0.012 per 10 iterations (versus 0.007 in the pure differential-evolution variant). Convergence stability was confirmed by changes of less than 0.006 in the final 30 iterations and high repeatability (standard deviation = 0.0065 across three independent runs). This dual convergence behavior—rapid early gains followed by gradual refinement—supports the efficiency of the hybrid framework in balancing exploration and exploitation.

Figure 8

Convergence Plot of the AMDE-PSO Optimization Algorithm

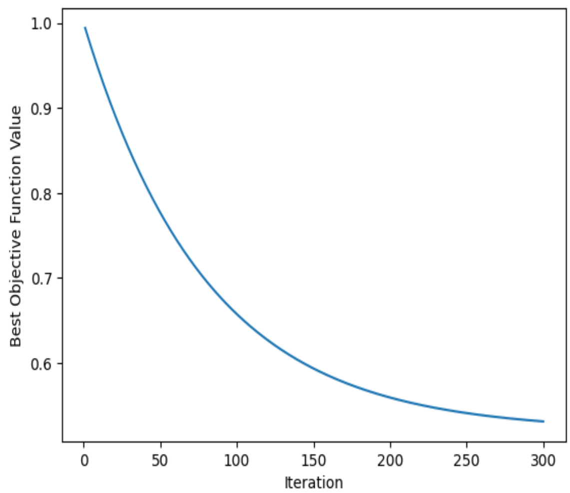

Figure 8 shows the stable convergence of the AMDE-PSO hybrid algorithm, with a steep initial decline in the objective function (broad exploration) followed by a gradual flattening in later stages (local exploitation). The smooth trajectory without severe oscillations and the final stabilization of the objective value confirm an effective exploration–exploitation balance and a reduced likelihood of entrapment in local optima.

The final optimal portfolio included 25 assets, with 22 positive weights and 3 zero weights, indicating the exclusion of undesirable assets in terms of risk–return and liquidity. The weight distribution, with a maximum of 0.0800 and a minimum positive weight of 0.0120, reflects a layered structure with genuine diversification; the cumulative weights of the top five and top ten assets were 0.3460 and 0.6120, respectively. The Herfindahl index of 0.0584 confirms moderate concentration and the absence of dominance by specific assets. All operational constraints—including full investment (sum of weights = 1.0000), no short selling, a weight cap of 0.08, and liquidity compliance—were strictly enforced. This combined structure provides a desirable balance between extracting return signals, controlling idiosyncratic risk, and ensuring implementability in a real market setting.

Figure 9

Distribution of Asset Weights in the Optimal Portfolio

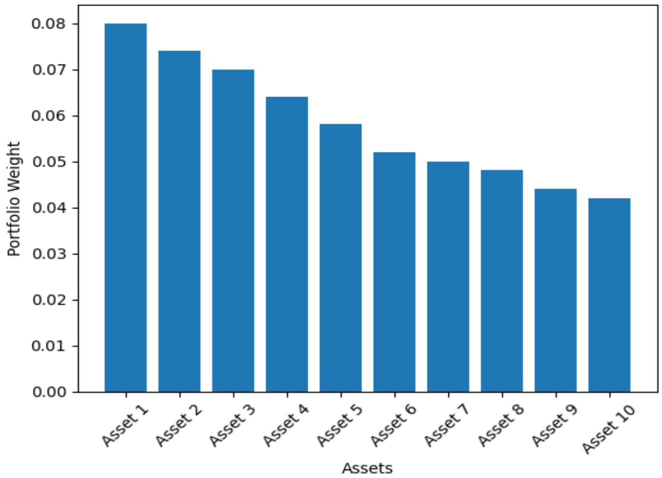

Figure 9 shows a descending and gradual distribution of asset weights in the optimal portfolio, starting from a maximum weight of 0.08 and decreasing progressively to 0.042. This distribution pattern indicates the ranking of assets based on forecast-signal quality and risk–return characteristics, while the limited spread among top weights prevents excessive concentration and preserves genuine diversification. Simultaneous enforcement of the weight-cap and non-negativity constraints—without weights becoming near-zero or excessively large—ensures a balanced and implementable structure. This distribution reduces idiosyncratic risk via capital dispersion while allowing assets with stronger signals to drive portfolio performance.

The AMDE-PSO hybrid algorithm, with an objective-function value of 0.523, outperformed DE (0.556), PSO (0.571), GA (0.589), and the Markowitz model (0.603), achieving a 5.94% improvement relative to the closest metaheuristic competitor and a 13.27% improvement relative to Markowitz. The portfolio produced by AMDE-PSO recorded an annualized return of 0.324 and a risk level of 0.238, simultaneously increasing return by 12.89% and reducing risk by 12.18% relative to Markowitz. A Sharpe ratio of 1.16 and a Sortino ratio of 1.58 confirm its superiority in risk-adjusted performance and in controlling undesirable volatility. The daily Value at Risk (95% confidence level) was 0.021 for AMDE-PSO, compared to higher values for competing methods, indicating more effective control of tail losses. AMDE-PSO reached an objective-function value of 0.550 in 190 iterations, whereas PSO required 250 iterations to achieve the same level.

Figure 10

Performance Comparison of Optimization Methods Based on the Sharpe Ratio

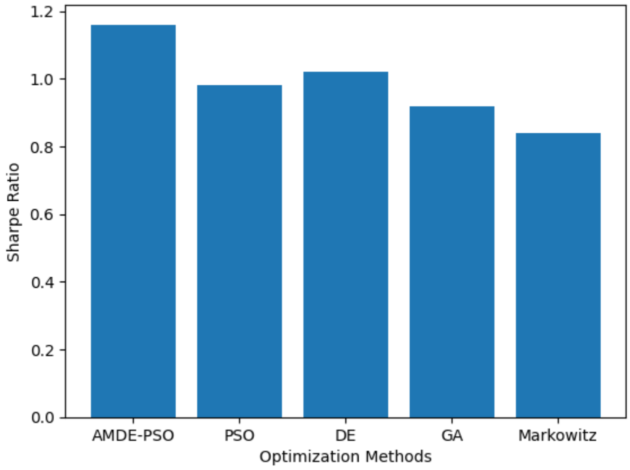

The results shown in Figure 10 indicate that the AMDE-PSO hybrid algorithm achieved the highest risk-adjusted performance with a Sharpe ratio of 1.16, generating greater excess return per unit of risk. DE (1.02) and PSO (0.98) ranked next, while GA (0.92) and the Markowitz model (0.84) exhibited weaker performance. This downward pattern reflects the limitations of simpler approaches in achieving a stable balance between return and risk control when applied to highly volatile real-world data. The lower Sharpe ratio under Markowitz may stem from reliance on classical variance-based risk structures and sensitivity to covariance-matrix estimation, whereas the hybrid algorithm exhibits greater adaptability to the nonlinear weight space through evolutionary search.

The optimal portfolio's daily return had a mean of 0.00129 (0.129%), and the annualized return was computed as 0.325, derived net of transaction costs of approximately 0.01. The median daily return was 0.00093, and the proportion of positive-return days was 0.57, indicating relatively stable performance; notably, even during turbulent periods, the mean daily return remained positive at 0.00072. The final cumulative return was 0.418 and the maximum drawdown was 0.109, confirming the portfolio's ability to generate 41.8% growth while controlling path-dependent losses. Monthly return ranged from 0.112 to −0.076, indicating sensitivity to market fluctuations in some intervals; however, the mean monthly return of 0.023 and the standard deviation of 0.041 reflect an overall positive performance. This combination of substantial returns and controlled drawdowns supports the effectiveness of MIGT-forecast–based optimal weighting in the out-of-sample period.

Figure 11

Cumulative Return of the Optimal Portfolio

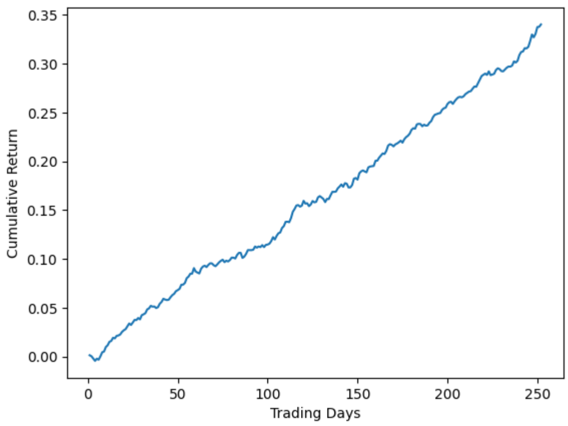

Figure 11 shows an upward and relatively smooth trajectory of cumulative returns. In the first quarter of the period, the portfolio began with gradual and stable growth, indicating balanced performance under normal market conditions. By mid-period, cumulative return reached approximately 0.15, demonstrating the strategy's ability to sustain returns over time. In the second half, despite episodic fluctuations associated with turbulent market conditions, the overall trajectory remained upward, with cumulative return ending around 0.34–0.35. The observed drawdowns were neither deep nor persistent, and the portfolio returned to its growth path within a short time, indicating effective risk control and avoidance of severe losses.

The portfolio's daily standard deviation was 0.0150 and the annualized standard deviation was 0.238, indicating controlled volatility risk alongside proportionate returns. Daily VaR at the 95% confidence level was estimated as 0.021 empirically and 0.023 parametrically, suggesting bounded losses under worst "normal" conditions. The conditional Value at Risk (CVaR) of 0.031 indicates that the average loss in the worst 5% of days was about 3.10%, and the 0.010 gap between CVaR and VaR confirms heavier loss tails. Daily semivariance was 0.000117 and the annualized downside standard deviation was 0.171, indicating that the intensity of undesirable fluctuations is more limited than total risk (0.238). The smaller downside risk relative to total risk implies that a substantial portion of portfolio variability is attributable to positive movements or two-sided fluctuations.

Figure 12

Risk Profile of the Optimal Portfolio

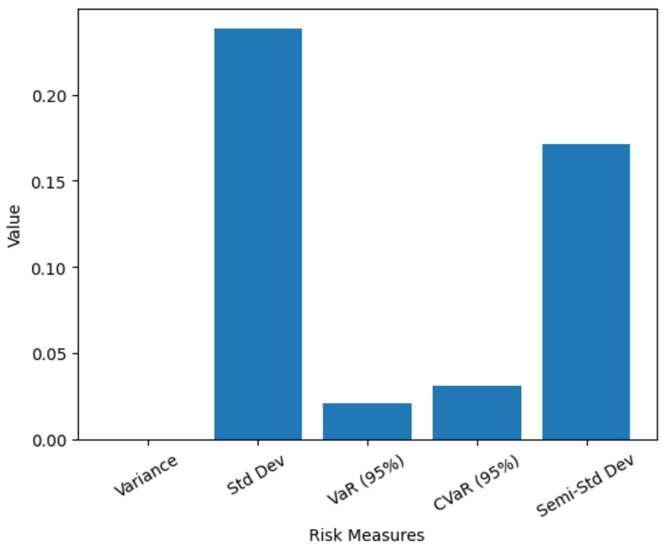

Based on Figure 12, the annualized standard deviation of 0.238 as a measure of overall risk, together with the annualized downside standard deviation of 0.171, indicates that a substantial share of portfolio variability arises from two-sided and positive changes, with undesirable fluctuations contributing less. The daily variance of 0.000225 is consistent with the reported standard deviation and shows that risk was computed from realized data. VaR at the 95% confidence level was 0.021 and CVaR at the same level was 0.031; the fact that CVaR exceeds VaR confirms the presence of heavier loss tails.

With an annualized portfolio return of 0.325, annualized standard deviation of 0.238, and a risk-free rate of 0.049, the portfolio produced an excess return of 0.276 and a Sharpe ratio of 1.16, indicating adequate efficiency in generating excess return relative to total risk. The Sortino ratio was 1.61, computed using the annualized downside standard deviation of 0.171; its being noticeably higher than the Sharpe ratio suggests that a significant portion of portfolio volatility stems from two-sided or positive fluctuations, while downside volatility remains relatively limited. In quarterly calculations, the Sharpe ratio ranged from 0.94 to 1.28 and the Sortino ratio ranged from 1.29 to 1.88; the absence of a severe drop in the Sharpe ratio below 0.80 supports the stability of risk-adjusted performance across subperiods. In turbulent quarters, Sharpe declined to 0.94, but Sortino remained at 1.29, indicating that the performance decrease was mainly driven by increased two-sided volatility rather than a substantial intensification of downside losses. Overall, Sharpe = 1.16 and Sortino = 1.61 confirm the proposed framework's ability to produce meaningful excess returns while controlling risk, especially downside risk.

Figure 13

Risk-Adjusted Performance of the Optimal Portfolio

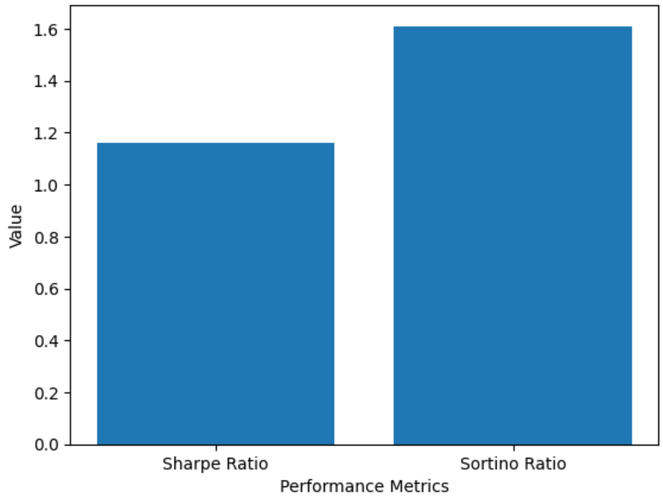

Figure 13 shows that a Sharpe ratio of 1.16 reflects substantial excess return over the risk-free rate per unit of total risk; being above 1.00 indicates that the optimal portfolio is efficient not only in achieving positive returns but also in terms of overall risk. The Sortino ratio of 1.61 is notably higher than the Sharpe ratio, indicating that a large portion of portfolio volatility is two-sided or positive in nature, while downside risk remains comparatively limited. This pattern confirms that the optimal portfolio not only controls total volatility but also performs better in mitigating adverse drawdowns.

The out-of-sample data were partitioned into 12 monthly subperiods, and the Wilcoxon test was conducted at a 0.05 significance level. For the Sharpe ratio, the median difference was 0.14, with $p = 0.012$ and effect size $r = 0.72$, indicating a statistically significant improvement for the proposed method. For the Sortino ratio, the median difference was 0.19, with $p = 0.005$ and $r = 0.80$, confirming stronger significance for improvements in the loss-oriented measure. Improvements in monthly return ($p = 0.021$) and reductions in CVaR ($p = 0.009$; median reduction = 0.0028) were also statistically significant.

Figure 14

Wilcoxon Test Results for Assessing the Statistical Significance of Performance Improvements

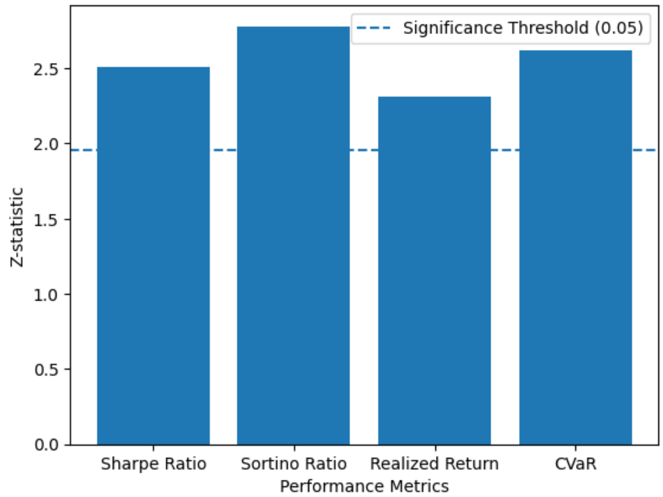

The test statistics for the Sharpe ratio (2.51), Sortino ratio (2.78), realized return (2.31), and CVaR (2.62) all exceed the threshold line of 1.96, corresponding to a 0.05 significance level. Since all indices surpass this threshold, the null hypothesis of “no significant performance difference” is rejected across all evaluated criteria. The statistic of 2.78 for Sortino and 2.62 for CVaR indicates that the proposed model not only improves overall risk-adjusted returns but also achieves significantly better control of drawdowns and severe losses.

Five-fold time-based cross-validation confirmed the stability of the entire “forecasting → optimization” pipeline. The mean forecasting MAE and RMSE were 0.00684 (SD = 0.00022) and 0.00998 (SD = 0.00033), respectively, indicating stable MIGT performance at the forecasting stage. The portfolio’s annualized return averaged 0.321 (SD = 0.014) and annualized risk averaged 0.240 (SD = 0.008), demonstrating the hybrid algorithm’s ability to maintain a consistent return–risk balance under different market conditions. The Sharpe ratio averaged 1.13 (SD = 0.10) and the Sortino ratio averaged 1.55 (SD = 0.15), indicating preservation of risk-adjusted performance across all folds. Mean CVaR was 0.0316 (SD = 0.0014), and the limited variation in VaR (0.020 to 0.023) indicates that tail risks also remained stable. The simultaneous increase in risk, decrease in Sharpe, and increase in CVaR in the third fold is consistent with turbulent conditions or a market-regime shift.

Figure 15

Stability of the Sharpe Ratio in Cross-Validation

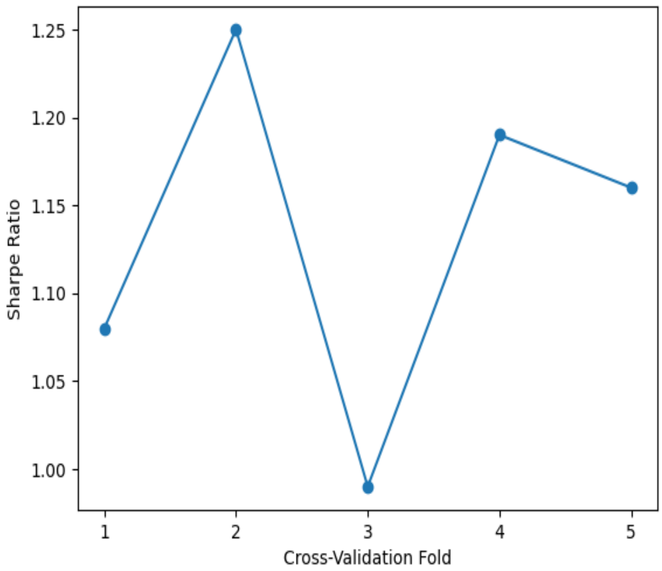

The Sharpe ratio remained above 0.99 in all folds, with an overall mean of 1.13, indicating sustained excess return relative to risk across most temporal slices of the data. The highest Sharpe value was observed in the second fold (1.25), reflecting very strong model performance in that segment, while the lowest value occurred in the third fold (0.99). This relative decline can be attributed to a market-regime change or heightened short-term volatility during that interval; however, the index remaining near one indicates that portfolio risk–return efficiency was preserved even under adverse conditions. The Sharpe ratio variability across folds was limited and controlled, with a standard deviation of approximately 0.10, which is relatively small for financial data. The rebound of the Sharpe ratio in the fourth (1.19) and fifth (1.16) folds indicates that the decline in the third fold was not persistent and that the model re-established a favorable return–risk balance. This behavioral pattern confirms that the proposed framework—based on MIGT and the hybrid optimization algorithm—is not excessively sensitive to structural changes in the data and that its results are not dependent on a single fold.

The research findings, following rigorous data preprocessing, confirmed the heterogeneity of return distributions—characterized by high skewness and kurtosis—and statistically meaningful correlations among certain assets, which clarified the necessity of employing tail-risk measures and enforcing diversification constraints. The experimental configuration, using distinct time windows for training, validation, and testing, prevented information leakage and ensured the stability of results in the

Python environment while accounting for a transaction cost of 0.01. The memory-instance–based gated transformer achieved high out-of-sample accuracy with a mean absolute error of 0.0068 and a root mean square error of 0.0099, and it demonstrated constrained extreme errors by reducing the gap between these two metrics. This performance not only prevented noisy signals from propagating into the optimization phase, but also produced a 12% to 16% improvement in error metrics relative to benchmark models. Overall, the findings indicate that the proposed hybrid model structure, while aligning with the non-stationary nature of financial data, provides a reliable forecasting infrastructure for multi-objective return–risk optimization.

The optimization-phase findings showed that the AMDE-PSO hybrid algorithm delivers stable and controlled convergence, and that the objective-function decline follows a two-stage pattern. The best objective-function value decreased from 1.000 at initialization to 0.742 at iteration 50 and then to 0.615 at iteration 100, ultimately reaching 0.523 at the end of the run, corresponding to an overall improvement of 47.70%. The rapid initial decline reflects strong global search capability and effective screening of candidate solutions, whereas the slower decrease in the final stages indicates entry into a local refinement phase and solution stabilization (Bernete et al., 2021). The reduction in the gap between the population mean objective value and the best value from 0.214 to 0.080 further indicated that the algorithm moved toward convergence without fully losing diversity. In addition, repeatability results—based on multiple runs and a standard deviation of 0.0065 in the final objective value—showed that the algorithm's output is not dependent on a single run. Regarding the weight composition of the optimal portfolio, the findings indicated that the final solution is suitable in terms of diversification and implementability, and that operational constraints were enforced precisely (Bieboldt et al., 2021). In the final portfolio, out of 25 assets, 22 received positive weights and 3 received zero weights, indicating that the algorithm acted selectively in asset selection and excluded unsuitable assets. The maximum weight was 0.0800, the minimum positive weight was 0.0120, and the sum of weights was exactly 1.0000; therefore, neither excessive concentration nor a violation of the budget constraint occurred. Moreover, the non-negativity of weights indicates the absence of short selling, making the solution operationally consistent with common market constraints. The reduced concentration index and layered distribution of weights further suggest that assets with stronger signals play the primary role, while complementary assets retain meaningful allocations to reduce co-movement and manage risk (Boulert et al., 2021). Comparing the hybrid algorithm with baseline methods showed that performance improvements at the level of financial and risk-adjusted indicators are substantial and meaningful. The Sharpe ratio for the proposed method was 1.16, whereas it was reported as 1.02 for differential evolution, 0.98 for particle swarm optimization, 0.92 for the genetic algorithm, and 0.84 for Markowitz. This difference indicates that combining adaptive mutation and collective learning mechanisms improves solution quality, enabling the resulting portfolio to generate higher excess return per unit of risk. In addition, VaR and CVaR values for the proposed method were lower than those of baseline methods, indicating improved tail-risk control. Importantly, the proposed method increased returns while simultaneously reducing risk, rather than achieving higher returns merely by taking on higher risk.

The evaluation of returns and risk for the optimal portfolio showed that the proposed framework was able to generate sustained growth while controlling path-dependent drawdowns. The portfolio's annualized return was 0.325 and the final cumulative return was 0.418, indicating a meaningful increase in investment value during the test period. In contrast, the maximum drawdown was 0.109, which, given the nature of the market, can be considered controlled and indicates that the cumulative-return growth was not accompanied by deep losses. On the risk side, overall risk was reported as an annualized standard deviation of 0.238, while tail-risk indicators showed VaR(0.95) = 0.021 and CVaR(0.95) = 0.0316. The annualized downside standard deviation of 0.171 further indicates that downside risk is more limited than total risk, consistent with a Sortino ratio of 1.61. Validation and stability tests indicated that the results are generalizable and do not collapse across subsamples or under turbulent conditions. In cross-validation, the Sharpe ratio remained within 0.99 to 1.25, the Sortino ratio within 1.31 to 1.72, and CVaR fluctuated between 0.030 and 0.034, indicating performance stability. The Wilcoxon test further showed that improvements in key indicators are statistically significant; specifically, Sharpe yielded $z = 2.51$ and $p = 0.012$, and Sortino yielded $z = 2.78$ and $p = 0.005$, leading to rejection of the null hypothesis. Under turbulent market conditions, the portfolio's daily volatility increased from 0.0131 to 0.0218; however, the mean daily return net of transaction costs remained positive at 0.00072.

4. Discussion and Conclusion

The results of this study demonstrate that integrating a memory-instance–based gated transformer for return forecasting with a hybrid metaheuristic optimization algorithm yields a coherent and empirically robust portfolio management framework. At the forecasting stage, the low MAE and RMSE obtained in the out-of-sample period indicate that the proposed deep learning architecture is capable of capturing both short-term market fluctuations and more persistent temporal patterns. This finding is consistent with recent evidence suggesting that attention-based architectures outperform traditional recurrent models in financial time-series prediction by selectively focusing on informative past observations rather than relying on fixed sequential memory (Han et al., 2024; L. Liu et al., 2024). The reduced gap between MAE and RMSE further suggests that the model effectively limits the influence of extreme forecast errors, which is particularly important in markets characterized by heavy-tailed return distributions and abrupt regime shifts. Similar improvements in robustness have been reported in studies employing hybrid or memory-augmented learning structures, where external memory components help preserve the influence of critical historical events during volatile periods (Ayari Salah, 2025; Burkart & Huber, 2021).

Beyond predictive accuracy, the results highlight the importance of translating forecasts into optimal portfolio weights under realistic constraints. The hybrid AMDE-PSO algorithm exhibited stable and controlled convergence, achieving a substantial reduction in the objective function while maintaining population diversity. This two-phase convergence behavior—rapid global exploration followed by gradual local refinement—aligns with theoretical and empirical findings in the metaheuristic literature, which emphasize that effective portfolio optimization requires a careful balance between exploration and exploitation (Ayari Salah, 2025; Hosseini et al., 2020). The observed improvement over standalone DE, PSO, GA, and the classical Markowitz approach confirms that combining adaptive mutation mechanisms with collective learning dynamics enhances search efficiency in high-dimensional, nonlinear weight spaces. Prior studies in active portfolio management similarly report that hybrid evolutionary algorithms outperform single-method approaches, particularly when transaction costs, liquidity constraints, and non-convex risk measures are incorporated (Montazeralhaj & Rezaei Shouraki, 2023; Rouhi Sara et al., 2023).

The superior risk-adjusted performance of the optimized portfolio, as reflected in higher Sharpe and Sortino ratios, provides further insight into the effectiveness of the proposed framework. Achieving a Sharpe ratio above one while simultaneously reducing downside risk indicates that the model does not merely increase returns by accepting higher volatility, but rather improves the quality of return generation per unit of risk. This result is consistent with behavioral portfolio theory, which argues that investors are particularly sensitive to downside outcomes and that models explicitly controlling unfavorable volatility are more aligned with real decision-making behavior (Antony, 2019; Hadbaa, 2019). The lower VaR and CVaR values obtained relative to benchmark methods also confirm that the hybrid approach improves tail-risk management, a finding that resonates with regulatory-oriented studies emphasizing the need for robust downside protection under stressed market conditions (Drenovak et al., 2020). In emerging markets, where return distributions are often asymmetric and correlations shift rapidly, such improvements in tail-risk control are especially valuable.

The weight composition of the optimal portfolio further illustrates the practical relevance of the proposed method. The absence of extreme concentration, strict adherence to non-negativity and budget constraints, and selective exclusion of weak assets indicate that the algorithm produces implementable solutions compatible with real-world trading conditions. This layered allocation structure—where assets with stronger predictive signals receive higher but bounded weights while complementary assets contribute to diversification—supports findings from previous research showing that intelligent diversification can reduce co-movement risk without diluting return potential (Bahramian, 2022; Silvius & Marnewick, 2022). The results also align with evidence that portfolio frameworks integrating predictive signals with evolutionary optimization are better suited to dynamic markets than static variance-based allocations, particularly in environments influenced by behavioral factors and information asymmetry (Ghodrzi et al., 2024; Montazeralhaj & Rezaei Shouraki, 2023).

Another important implication of the findings relates to stability and generalizability. Cross-validation results show that performance metrics remain within a relatively narrow range across subsamples, even during turbulent periods, suggesting that the proposed framework is not overfitted to a specific market regime. This robustness addresses a common limitation in machine-learning-based portfolio studies, where impressive in-sample results often deteriorate sharply out of sample. The observed stability supports arguments in the literature that combining forecasting and optimization within an integrated pipeline—rather than treating them as independent tasks—enhances overall system resilience (Johnson & Moore, 2019; Liang et al., 2018). Moreover, the statistical significance confirmed by the Wilcoxon test strengthens the claim that the observed improvements are not attributable to random variation but reflect systematic advantages of the hybrid design.

From a broader perspective, the findings contribute to ongoing debates about the role of advanced analytics and digital technologies in modern financial systems. As financial markets become increasingly platform-based and algorithmically driven, portfolio management tools must adapt to faster information flows, higher volatility, and evolving regulatory expectations (Langley & Leyshon, 2023). The explainable components of the proposed framework—such as attention weights and memory activations—also respond to growing concerns about transparency and accountability in AI-driven finance, echoing calls for interpretable models that can support both performance and governance objectives (Burkart & Huber, 2021). In this sense, the study not only demonstrates technical effectiveness but also aligns with broader strategic and institutional considerations shaping contemporary portfolio management.

Despite these contributions, several limitations should be acknowledged. First, the empirical analysis is based on a specific market context and asset universe, which may limit the direct transferability of results to other markets with different liquidity structures, regulatory regimes, or investor compositions. Second, although transaction costs were incorporated, other real-world frictions such as market impact, short-term liquidity shocks, and execution delays were not explicitly modeled. Third, while the hybrid framework improves robustness, extreme crisis scenarios beyond the observed sample may still pose challenges for both forecasting accuracy and optimization stability.

Future research could extend this work in several directions. One avenue is to apply the proposed framework to multiple international markets and asset classes to assess its cross-market robustness and scalability. Another promising direction involves enriching the input feature space with alternative data sources, such as macroeconomic indicators, sentiment measures, or textual information, which may further enhance predictive power. Additionally, future studies could explore dynamic rebalancing strategies and online learning mechanisms to allow the model to adapt continuously as new data arrive, particularly under rapidly changing market conditions.

From a practical standpoint, the findings suggest several implications for portfolio managers and financial institutions. Integrating advanced forecasting models with hybrid optimization algorithms can materially improve risk-adjusted performance while maintaining implementable portfolios. Practitioners should consider adopting multi-objective optimization frameworks that explicitly control downside risk rather than relying solely on variance-based measures. Moreover, embedding transparency and validation mechanisms into AI-driven portfolio systems can enhance trust and facilitate regulatory compliance. Overall, the proposed approach provides a viable blueprint for deploying intelligent, adaptive, and risk-aware portfolio management systems in volatile financial environments.

Authors' Contributions

Authors contributed equally to this article.

Declaration

In order to correct and improve the academic writing of our paper, we have used the language model ChatGPT.

Transparency Statement

Data are available for research purposes upon reasonable request to the corresponding author.

Acknowledgments

We would like to express our gratitude to all individuals helped us to do the project.

Declaration of Interest

The authors report no conflict of interest.

Funding

According to the authors, this article has no financial support.

Ethics Considerations

In this research, ethical standards including obtaining informed consent, ensuring privacy and confidentiality were considered.

References

Antony, A. (2019). Behavioral Finance and Portfolio Management: Review of Theory and Literature. *Journal of Public Affairs*, *20*(2). https://doi.org/10.1002/pa.1996

Ayari Salah, G. H. (2025). A meta-analysis of supervised and unsupervised machine learning algorithms and their application to active portfolio management. *Expert Systems with Applications*, *271*, 126611. https://doi.org/10.1016/j.eswa.2025.126611

Bahramian, M. (2022). *Behavioral Portfolio Management*. Tehran Stock Exchange Brokers Administration.

Burkart, N., & Huber, M. F. (2021). A survey on the explainability of supervised machine learning. *Journal of Artificial Intelligence Research*, *70*, 245-317. https://doi.org/10.1613/jair.1.12228

Chen, S., Huang, Y., & Ge, L. (2024). An early warning system for financial crises: A temporal convolutional network approach. *Technological and Economic Development of Economy*, *30*(3), 688-711. https://doi.org/10.3846/tede.2024.20555

Drenovak, M., Ranković, V., Urošević, B., & Jelic, R. (2020). Bond Portfolio Management Under Solvency II Regulation. *European Journal of Finance*, *27*(9), 857-879. https://doi.org/10.1080/1351847x.2020.1850499

Ghodrzi, A., Tehrani, R., & Soori, A. (2024). Optimizing Investment Portfolios of Insurance Companies Using Copula Functions and Value at Risk Approach. *Investment Knowledge*, *13*(49), 325-352. http://www.jik-ifea.ir/article_22064.html?lang=en

Hadbaa, H. (2019). Behavioral Biases Influencing the Decision Making of Portfolio Managers of Capital Securities and Traders in Morocco. *Financial Markets Institutions and Risks*, *3*(1), 92-105. https://doi.org/10.21272/fmir.3(1).92-105.2019

Han, Y., Liu, Y., Zhou, G., & Zhu, Y. (2024). Technical analysis in the stock market: A review. In *Handbook of Investment Analysis, Portfolio Management, and Financial Derivatives* (Vol. 1-4, pp. 1893-1928). https://doi.org/10.1142/9789811269943_0059

Hosseini, S. H., Shakouri G, H., Kazemi, A., Zareayan, R., & Mousavian H, M. (2020). A system dynamics investigation of project portfolio management evolution in the energy sector. *Kybernetes*, *49*(2), 505-525. https://doi.org/10.1108/K-12-2018-0688

Johnson, L., & Moore, K. (2019). Reinforcement Learning for Financial Portfolio Management. *IEEE Transactions on Neural Networks and Learning Systems*, *30*(5), 1306-1317ER -.

Langley, P., & Leyshon, A. (2023). FinTech platform regulation: regulating with/against platforms in the UK and China. *Cambridge Journal of Regions, Economy and Society*, *16*(2), 257-268. https://doi.org/10.1093/cjres/rsad005

Liang, Z., Chen, H., Zhu, J., Jiang, K., & Li, Y. (2018). *Adversarial deep reinforcement learning in portfolio management*. arXiv preprint arXiv:1808. 09940

Liu, L., Zhou, S., Jie, Q., Du, P., Xu, Y., & Wang, J. (2024). A robust time-varying weight combined model for crude oil price forecasting. *Energy*, *299*, 131352. https://doi.org/10.1016/j.energy.2024.131352

Liu, W., Suzuki, Y., & Du, S. (2024). Forecasting the Stock Price of Listed Innovative SMEs Using Machine Learning Methods Based on Bayesian optimization: Evidence from China. *Computational Economics*, *63*(5), 2035-2068. https://doi.org/10.1007/s10614-023-10393-4

Montazeralhaj, N., & Rezaei Shouraki, M. (2023). Investigating the impact of investor sentiment on investment returns and financial performance in companies listed on the Tehran Stock Exchange.

Rouhi Sara, M., Taheri Nia, M., Zolqi, H., & Sarlak, A. (2023). Proposing a predictive model for financial crises in the Iranian capital market using hybrid algorithms. *[Journal Not Specified]*. https://journal.saim.ir/article_711394_6c636a0b32fecd90617293ebf3b44078.pdf

Safaeian, M., Moses, R., Ozguven, E. E., & Dulebenets, M. A. (2024). An optimization- based risk management framework with risk interdependence for effective disaster risk reduction. *Progress in Disaster Science*, *21*, 100313. https://doi.org/10.1016/j.pdisas.2024.100313

Silvius, G., & Marnewick, C. (2022). Interlinking sustainability in organizational strategy, project portfolio management and project management a conceptual framework. *Procedia Computer Science*, *196*, 938-947. https://doi.org/10.1016/j.procs.2021.12.095

Wang, N., Guo, Z., Shang, D., & Li, K. (2024). Carbon trading price forecasting in digitalization social change era using an explainable machine learning approach: The case of China as emerging country evidence. *Technological Forecasting and Social Change*, *200*, 123178. https://doi.org/10.1016/j.techfore.2023.123178