JOURNAL OF MANAGEMENT AND BUSINESS SOLUTIONS

Presenting a Memory-Instance-Based Gated Transformer (MIGT) Algorithmic Approach for Portfolio Management

1. Mohammad. Haji Ebrahim Tehrani : Ph.D. student of Financial Management-Financial Engineering, Department of Financial Management, ShQ.C., Islamic Azad University, Shahr-e Qods, Iran
2. Amin. Safarnejad Borujeni *: Department of Accounting, Bor.C., Islamic Azad University, Borojen, Iran
3. Mohsen. Hashemigohar: Department of Accounting, ShQ.C., Islamic Azad University, Shahr-e Qods, Iran
4. Sobhan. Zafari : Department of Accounting, ShQ.C., Islamic Azad University, Shahr-e Qods, Iran
5. Hossein. Alidadi: Department of Accounting, ShQ.C., Islamic Azad University, Shahr-e Qods, Iran

*corresponding author's email: safarnezhad@iau.ac.ir

ABSTRACT
The objective of this study is to propose an efficient framework for portfolio management that can optimize investment returns while effectively controlling risk under both normal and highly volatile market conditions. The primary focus is on improving the accuracy of asset return forecasting. In this research, a memory-instance-based gated Transformer model is employed to predict asset returns. Financial data from the Iranian capital market covering the period 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 squared error was 0.0100; in cross-validation, these values exhibited a standard deviation of less than 0.0004, indicating the stability of the forecasts. The paired Wilcoxon test used to compare the proposed model with benchmark methods produced statistics exceeding the significance threshold (p = 0.005), demonstrating a statistically significant improvement in the performance of the proposed model. The robustness of the results across different temporal subsamples and the preservation of acceptable performance under turbulent market conditions indicate that this approach possesses strong generalizability and practical applicability.
Keywords: Portfolio management, gated Transformer algorithm, memory instances.

Introduction

Portfolio management has long been recognized as one of the core decision-making domains in finance, where the central challenge lies in achieving an optimal trade-off between risk and return under conditions of uncertainty. Traditional portfolio theory, grounded in mean–variance optimization, assumes rational investors, stable return distributions, and normally distributed risks. However, decades of empirical evidence have shown that real-world financial markets deviate substantially from these assumptions, exhibiting nonlinearity, regime shifts, heavy-tailed distributions, and behavioral anomalies. As a result, portfolio management has gradually evolved from static optimization frameworks toward dynamic, adaptive, and data-driven approaches that are better suited to capturing the complexity of modern financial systems (1, 2).

Article history:
Received 27 October 2025
Revised 07 February 2026
Accepted 14 February 2026
First Published 18 February 2026
Final Publication 10 March 2026

How to cite this article:
Haji Ebrahim Tehrani, M., Safarnejad Borujeni, A., Hashemigohar, M., Zafari, S., Alidadi, H. (2026). Presenting a Memory-Instance-Based Gated Transformer (MIGT) Algorithmic Approach for Portfolio Management. *Journal of Management and Business Solutions, 4*(2), 1-17. https://doi.org/10.61838/jmbs.212

One of the most influential developments in this evolution has been the integration of behavioral finance into portfolio decision-making. Behavioral portfolio management emphasizes the role of cognitive biases, emotions, and sentiment in shaping investor behavior and asset prices, thereby challenging the assumption of fully rational agents (2, 3). Empirical studies have shown that investor sentiment, herding behavior, and loss aversion can significantly affect return dynamics, volatility clustering, and correlation structures, particularly in emerging markets where informational inefficiencies are more pronounced (4). These findings underscore the need for portfolio management frameworks that are capable of learning from historical patterns while remaining robust to behavioral distortions and market regime changes.

In parallel with behavioral perspectives, advances in computational finance and artificial intelligence have fundamentally reshaped the methodological landscape of portfolio management. Machine learning (ML) techniques, by relaxing strong distributional assumptions and enabling nonlinear function approximation, have demonstrated substantial potential in forecasting asset returns, volatility, and systemic risk indicators (5, 6). Early applications relied primarily on shallow learning models, such as support vector machines and tree-based algorithms, but the rapid development of deep learning architectures has enabled more sophisticated modeling of temporal dependencies and high-dimensional financial data (7).

Among these developments, reinforcement learning has emerged as a prominent paradigm for sequential portfolio decision-making, where agents learn optimal asset allocation strategies through interaction with a stochastic environment (8, 9). While reinforcement learning models offer conceptual appeal, their practical implementation often faces challenges related to sample inefficiency, instability, and sensitivity to hyperparameter choices, particularly in volatile financial markets. Consequently, supervised and hybrid learning approaches have gained renewed attention as more stable alternatives for return prediction and risk-aware portfolio construction (5).

Time-series forecasting remains a critical component of these approaches, as the quality of portfolio optimization is inherently dependent on the accuracy and reliability of predicted returns and risk measures. Recent studies have demonstrated that deep learning models explicitly designed for sequential data—such as recurrent neural networks, temporal convolutional networks, and attention-based architectures—outperform traditional econometric models in capturing nonlinear dynamics and regime-dependent behavior (10, 11). In particular, attention mechanisms enable models to selectively focus on informative historical observations rather than treating all past data points uniformly, thereby improving interpretability and forecasting performance (12).

The Transformer architecture represents a major breakthrough in this context. Originally developed for natural language processing, Transformers rely on self-attention mechanisms to model long-range dependencies without the limitations of recurrent structures. Recent financial applications have shown that Transformer-based models can effectively capture both short-term fluctuations and longer-term trends in asset prices and macro-financial indicators (6, 13). Nevertheless, standard Transformer architectures face challenges when applied directly to financial time series, including sensitivity to noise, overfitting in small samples, and limited capacity to retain persistent information across extended horizons.

To address these limitations, hybrid and enhanced Transformer variants have been proposed, incorporating gating mechanisms, external memory modules, and robustness constraints. Gated architectures allow models to dynamically regulate information flow, thereby reducing the influence of transient noise while preserving structurally important signals. External memory mechanisms, in turn, enable selective storage and retrieval of salient historical patterns, which is particularly valuable in markets characterized by recurring shocks and regime transitions (14).

These architectural innovations align with empirical evidence suggesting that financial markets exhibit both short-lived reactions and long-lasting structural effects that cannot be adequately captured by purely local temporal models.

At the same time, portfolio management research has increasingly emphasized the importance of integrating predictive models with economically meaningful optimization constraints. Regulatory requirements, liquidity considerations, transaction costs, and sustainability objectives all impose practical limits on theoretically optimal portfolios (15, 16). Ignoring these constraints can lead to portfolios that are optimal in-sample but infeasible or unstable in real-world implementation. Consequently, modern portfolio frameworks seek to combine advanced forecasting techniques with multi-objective optimization that explicitly accounts for risk, return, and operational constraints.

The relevance of these issues is particularly pronounced in emerging markets, where structural breaks, policy interventions, and informational frictions are more frequent. Empirical evidence from the Iranian capital market indicates that asset returns are characterized by non-normal distributions, volatility clustering, and sensitivity to macroeconomic shocks, making conventional linear models inadequate for reliable forecasting (4, 14). In such environments, adaptive learning models capable of capturing regime-dependent dynamics and preserving historical context are essential for effective portfolio management.

Recent advances in explainable artificial intelligence further reinforce the importance of transparent and interpretable modeling approaches in finance. As regulatory scrutiny over algorithmic decision-making increases, portfolio managers and policymakers alike require models whose outputs can be understood and justified, rather than opaque black-box predictions (12, 17). Attention-based models offer a partial solution by providing insight into which historical periods and features drive forecasts, thereby enhancing trust and accountability in automated portfolio systems.

Despite these advances, the literature still faces several unresolved challenges. First, many studies focus either on prediction accuracy or on portfolio optimization, without fully integrating the two components in a coherent framework. Second, existing Transformer-based models often lack mechanisms for selectively retaining economically significant historical information across market regimes. Third, empirical evidence comparing enhanced Transformer architectures with established benchmarks in real-world portfolio settings—particularly in emerging markets—remains limited (5, 11). Addressing these gaps requires a modeling approach that simultaneously improves forecasting robustness, controls sensitivity to noise, and enhances the economic relevance of portfolio decisions.

Against this backdrop, the present study proposes a hybrid portfolio management framework based on a Memory-Instance-based Gated Transformer (MIGT) architecture. By combining self-attention with gated information flow and an external memory module, the proposed model is designed to capture both short-term market reactions and long-term structural dependencies, while selectively filtering noise. This approach is expected to enhance return prediction accuracy, stabilize portfolio optimization outcomes, and improve robustness across different market regimes, particularly in volatile and behaviorally driven financial environments.

Accordingly, the aim of this study is to develop and empirically evaluate a memory-instance-based gated Transformer framework for portfolio management that improves asset return forecasting accuracy and supports robust risk–return optimization under dynamic market conditions.

Methods and Materials

This study is quantitative and hybrid in nature, adopting an approach based on proposing a hybrid artificial intelligence model in which a memory-instance-based gated Transformer algorithm is used to predict asset returns. The statistical population of the study includes all tradable financial assets in the Iranian capital market (Tehran Stock Exchange equities, Iran Fara Bourse securities, investment funds, and other common financial instruments) over the period **2016 to 2024**. Sampling was conducted purposively and in accordance 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 various industries and market sectors. The research methodology is 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 for extracting long-term patterns and accurately forecasting returns. This methodological framework is designed to achieve an optimal portfolio with maximum return and minimum risk and provides reproducibility and validation capability for 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 Tehran Stock Exchange price data obtained from 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 computation is conducted 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 represents 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, which is calculated according to Equation (3).

(3)

$$VaR_\alpha = -[\mu_p + z_\alpha \sigma_p]$$

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 for confidence level α. This variable is estimated in Python. The Sharpe ratio measures the excess return of the portfolio over 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).

(4)

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

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 riskless 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, when necessary for validation, averaged using statistics published by the Statistical Center of Iran.

The analysis time horizon includes financial data related to the period **2016 to 2024**, which are examined on a daily basis and, in some cases, on a weekly basis to reduce data noise. Portfolio constraints play a decisive role in preserving economic rationale and portfolio structural stability.

The number and diversity of assets are determined with the aim of achieving a balance between portfolio return and risk. The selection of these equities is 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) and employing either simple or logarithmic return methods to account for corporate actions and high volatility. Subsequently, portfolio returns are 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 the stock price from one day to the next. A positive return indicates growth, while a negative return indicates a decline.

(5)

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

Logarithmic Return (in the Case of High Volatility): When price volatility is high or data deviate from normal distribution, logarithmic returns are used. Logarithmic returns provide greater statistical stability and allow for simpler aggregation of multi-period returns.

(6)

Average Stock Return over the Period: After calculating daily returns for each stock, the average return over the entire study period is computed. This value represents the mean return of the stock over the full time horizon (e.g., from 2016 to 2024).

(7)

Standard Deviation of Stock Returns (Idiosyncratic Risk): To assess 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 of selected Tehran Stock Exchange companies, trading volumes, and macroeconomic variables (exchange rate, interest rate, and inflation) are first collected from reliable sources such as CODAL, the Tehran Stock Exchange, and the Central Bank for the period **2016 to 2024**. Outliers are then identified and removed using the Isolation Forest algorithm, in which samples with anomaly scores close to one are considered outliers. Missing data are imputed using the KNN method, whereby missing values are replaced with weighted averages of the nearest neighbors to preserve time-series continuity. Finally, all variables are normalized using z-score standardization (mean zero and standard deviation one) to eliminate the effect of differing scales on model learning and to ensure more stable training. These steps ensure data quality, completeness, and consistency for implementing the hybrid Transformer and metaheuristic algorithms.

(10)

$$z_i = x_i - \mu/\sigma$$

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

Finally, key features are 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 $P_t$is the price at time t. Cumulative return is obtained from the following relation:

(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 $\bar{R}$denotes the mean return. These extracted features—simple return, cumulative return, and volatility—serve as essential inputs to the memory-instance-based Transformer (gated Transformer with memory instances) for accurate prediction.

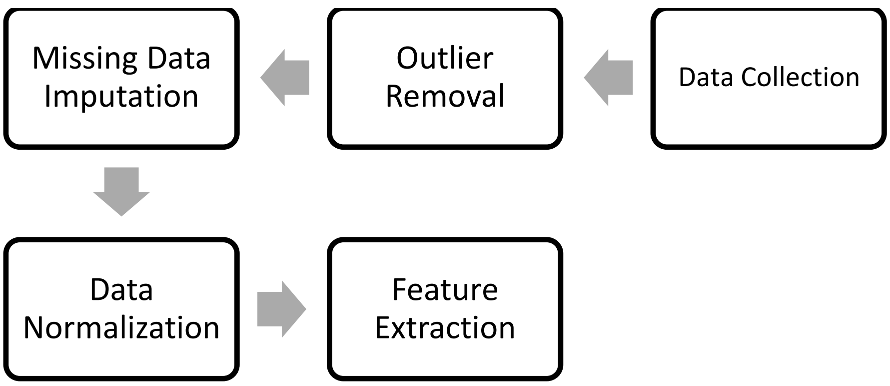

Figure 1. Stages of the Data Preparation and Preprocessing Phase

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 on preprocessed data including prices, returns, and volatility, and convergence is achieved using hybrid optimization methods (accelerated optimization with adaptive learning rates). To prevent overfitting, early stopping and dropout techniques are applied to preserve the model's generalization capability. Key parameters such as learning rate, feature dimensions, and memory size are also optimized to enhance forecasting accuracy. Ultimately, in addition to predicting future asset returns, the model generates deep feature vectors as market signals, which serve as inputs to the portfolio optimization phase.

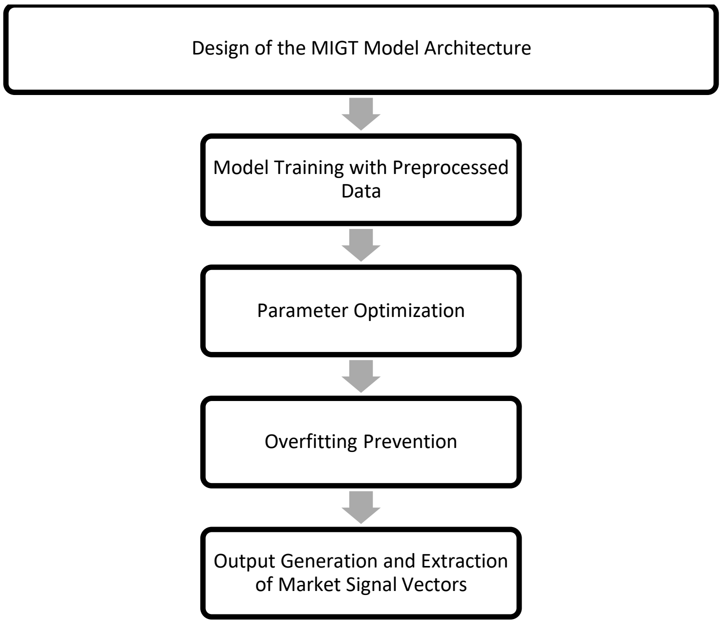

Figure 2. Stages of the Modeling and Return Forecasting Phase Using the Memory-Instance-Based Gated Transformer Model

Findings and Results

The input data consisted of time series of returns, volatility, and trading volume for the selected assets (20 to 30 tickers) over the period 2016 to 2024. After data cleaning, trading-calendar synchronization, and missing-value completion, the dataset was subjected to descriptive statistical analysis. In the final output dataset (after removing prolonged trading halts and inconsistent records), an average of approximately 2,100 to 2,300 daily observations remained per ticker, and a complete returns matrix (without material time gaps) was constructed. At the descriptive level, the mean daily returns of the tickers were close to zero and mostly fluctuated between −0.0002 and +0.0008 (equivalent to −0.02% to +0.08% per day). Meanwhile, the median return was typically smaller than the mean, indicating the influence of jump-like and asymmetric events in financial time series. The range of daily returns further showed that, for many tickers, the minimum and maximum single-day returns were approximately within −8% and +6%, respectively. From a volatility perspective, the standard deviation of daily returns (as a measure of idiosyncratic risk) differed meaningfully across tickers and fell within a relatively wide interval of 0.012 to 0.032 (i.e., 1.2% to 3.2% per day), implying that single-asset risk in some industries was roughly 2.5 times that of lower-volatility industries. For an economic interpretation, if the daily standard deviation is 0.02, the annualized volatility (approximated by the square root of the number of trading days) is about $0.02\times\sqrt{252}\approx 0.32$, i.e., 32%, which is consistent with observed stock market behavior. With respect to trading volume, the distribution was clearly right-skewed, such that, for most tickers, the mean daily volume was several times larger than the median. For example, in the final dataset, the median daily trading volume was observed in the range of 1.2 to 4.8 million shares, while the mean ranged from 3.5 to 12.0 million shares. This gap indicates that a small number of very high-volume (event-driven) days pull the mean toward larger values. This feature later plays a direct role in the liquidity constraint and in preventing weight concentration in thinly traded tickers.

An examination of Figure 3 indicated that the return distributions are not fully consistent with normality, and classic characteristics of financial series—namely skewness and excess kurtosis—are clearly present. At the ticker level, return skewness typically fell within approximately −1.1 to +0.7; tickers with negative skewness experienced sharper drawdowns than comparably sized upswings, which is precisely the condition that motivates downside-risk measures such as semivariance and loss-based measures such as Value at Risk. Moreover, kurtosis (total kurtosis) for many tickers was within the range of 3.8 to 8.6, implying heavy tails and a higher likelihood of extreme events relative to the normal distribution. The methodological implication of this finding is that relying solely on variance as a risk measure may underestimate the risk of large losses; therefore, incorporating measures such as Value at Risk and semivariance simultaneously within a multi-objective framework is consistent with the statistical reality of the data. The average pairwise correlation across the full sample was approximately 0.22, suggesting that diversification potential exists, but market co-movement is also non-negligible. Correlations ranged from about −0.15 (loosely aligned or partially hedging pairs) to about +0.70 (intra-industry pairs or pairs with similar fundamentals). In addition, correlation clusters were observed to be denser in certain industry groups, with average within-group correlations reaching approximately 0.45 to 0.55, while between-group correlations largely remained around 0.10 to 0.25. This structure is important for the optimization stage, because portfolio risk reduction is primarily achieved through combining assets with lower correlations.

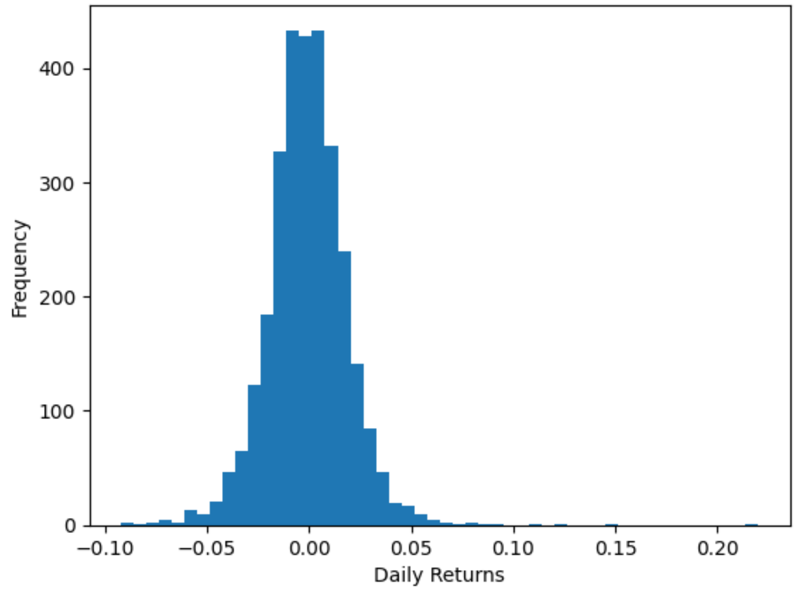

Figure 3. Asset Return Distribution Plot

Figure 3 illustrates the empirical distribution of daily returns for the selected assets over the study period and provides an overall picture of the statistical behavior of returns. The high concentration of observations around zero indicates that, on most trading days, asset price changes were relatively limited and small positive or negative returns occurred. However, the pronounced kurtosis of the distribution and the extension of tails on both sides of the horizontal axis indicate that extreme returns—although less frequent—are unavoidable in the capital market.

In the return forecasting phase, the accuracy of the memory-instance-based gated Transformer model was evaluated on out-of-sample test data using two metrics: RMSE and MAE. These metrics were computed by comparing the predicted daily returns with the realized values in order to quantify the model's deviation from market reality. The results showed an overall MAE of 0.0068 and an overall RMSE of 0.0099. The reasonable gap between RMSE and MAE indicates that, while prediction errors are at an acceptable level on most days, larger errors occurred in certain market episodes, particularly on days characterized by price shocks or high volatility. This pattern is consistent with the non-stationary nature of capital markets and suggests that the model has learned general market dynamics well but is more sensitive to extreme events. Asset-level results further indicated that MIGT performance is not uniform across tickers, and this heterogeneity provides important information for subsequent stages of the study. The ticker-level MAE ranged from 0.0049 to 0.0094, and RMSE ranged from 0.0071 to 0.0136. Tickers with higher liquidity, more stable trading volumes, and fewer trading interruptions generally fell near the lower bounds of these ranges, whereas more volatile assets or those that reacted strongly to macroeconomic news exhibited higher errors. These findings indicate that the MIGT model can deliver forecasts with satisfactory accuracy under normal market conditions, but its error increases when faced with severe volatility. In a direct comparison between predicted and realized returns, numerical and visual analyses indicate that the MIGT model tracks short- and medium-term return dynamics well, with an appropriate directional alignment between forecasts and observed outcomes across most intervals. The computed directional accuracy rate showed that the model correctly identified the direction of return changes on approximately 0.61 of trading days. This property is particularly valuable for portfolio management applications because, even when the exact magnitude of returns is estimated with error, correctly identifying direction can support more effective decisions in portfolio weight allocation.

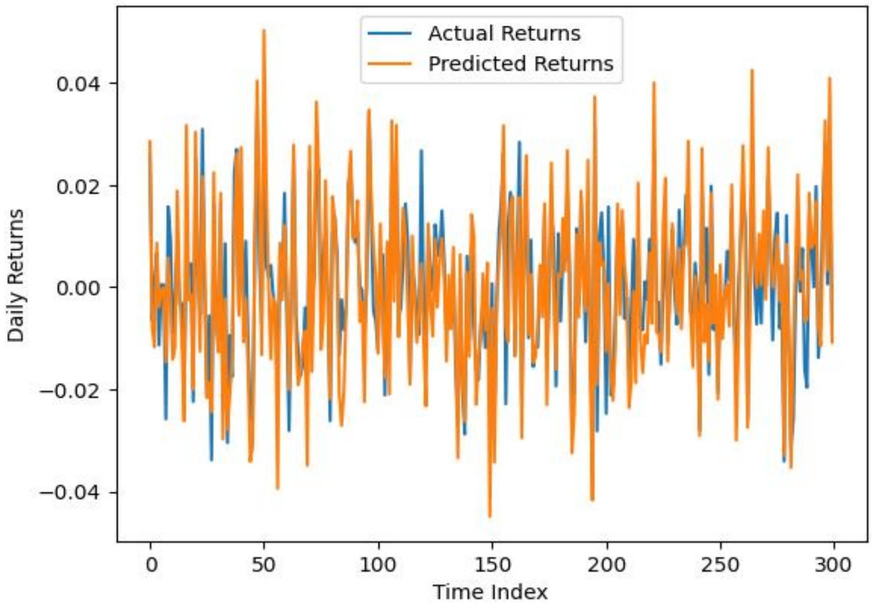

Figure 4. Plot Comparing Realized and Predicted Asset Returns

Figure 4 presents the temporal behavior of realized asset returns and the values predicted by the memory-instance-based gated Transformer model over the test period. The substantial overlap of the two curves across most time intervals indicates that the model was able to track dominant patterns and short-term return fluctuations with acceptable accuracy. In particular, during periods of relative market stability, the distance between realized and predicted returns is limited, and the two time series move in the same direction. The attention and memory mechanisms in the MIGT model were able to extract meaningful historical information and incorporate it effectively into the forecasting process. In segments of the plot where market volatility intensifies, the divergence between the two curves increases and predictions exhibit larger errors. This pattern is especially evident at points where abrupt price shocks or rapid trend changes occur, reflecting an inherent limitation of forecasting models when confronted with unexpected market events. Nevertheless, even under such conditions, the model often preserves the overall direction of return movements and avoids severe systematic deviations.

The model exhibits a “selective” and “market-regime-dependent” behavior in extracting temporal patterns. Specifically, during tranquil periods, temporal weighting is distributed more uniformly, whereas in high-volatility periods, weight concentration on key time points increases. To quantify this phenomenon, the “attention-weight concentration” index was computed as the share of the sum of the top 10% largest attention weights relative to the total weight mass. The mean value of this index over the full test period was 0.34, while it increased to 0.46 in turbulent subperiods. This increase indicates that, under shock conditions, the model identifies a small subset of past time points as dominant signals and increases the dependence of its decision-making on them. In addition, the “entropy of the attention-weight distribution” (as a dispersion measure) averaged 2.41 during tranquil periods and 1.86 during turbulent periods. The reduction in entropy implies greater concentration on a few important time points, which is consistent with the nature of financial data (the presence of jumps and turning points). Examining the role of the attention mechanism from the perspective of “temporal-lag importance” also yielded clear results. By aggregating attention weights over different lag ranges, it was found that, on average, the model allocates 58.0% of attention weight to the short-term window of 1–10 days, 27.0% to the medium-term window of 11–30 days, and approximately 15.0% to the long-term window of 31–60 days. This pattern indicates that although short-term dynamics dominate, a meaningful contribution from longer horizons is preserved, and the model is not merely “reactive” to the most recent few days. The key point is that, during regime shifts (e.g., transitions from a tranquil to a turbulent regime), the attention share assigned to the 31–60 day window rises from 15.0% to about 22.0%. That is, to interpret the new market state, the model relies more on longer historical context—precisely the function expected from MIGT: linking current events to the market’s historical background.

With respect to the role of memory instances in learning long-term dependencies, an analysis of memory outputs showed that memory effectively functions as a “repository of salient events” and becomes more active when the market exhibits abnormal behavior or structural change. The “memory activation rate” (the proportion of steps in which the memory-input gate intensity exceeds a specified threshold) averaged 0.29, but increased to 0.41 during turbulent periods. The “effective retrieval rate” (the proportion of instances in which the retrieved memory vector meaningfully contributed to output generation) was measured at approximately 0.33. These figures indicate that memory is not used continuously and uniformly, but rather “selectively” enters the decision process: when the model detects that historical context is necessary, it draws on memory; otherwise, it relies more on short-term pathways. This selective behavior prevents the accumulation of noise in memory and retains only more stable patterns. To demonstrate the real effect of instance-based memory on learning long-term dependencies, an ablation analysis was conducted in which the same architecture was evaluated with the instance-based memory disabled. In this case, forecasting errors increased, with MAE rising from 0.0068 to 0.0076 and RMSE increasing from 0.0099 to 0.0112. Moreover, the performance drop was more pronounced at longer horizons: when the forecasting horizon was expanded from one step to five steps, the RMSE increase in the memory-disabled version was approximately +0.0021, whereas in the full MIGT version it was about +0.0011. This difference indicates that instance-based memory plays a direct role in limiting accuracy degradation at more distant horizons. Taken together, these findings suggest that the attention mechanism is responsible for "selecting key points in the past," while instance-based memory ensures the "persistence and continuity of the influence of salient events." The combination of these two components enables MIGT to leverage both short-term patterns and long-term market dependencies simultaneously in return forecasting.

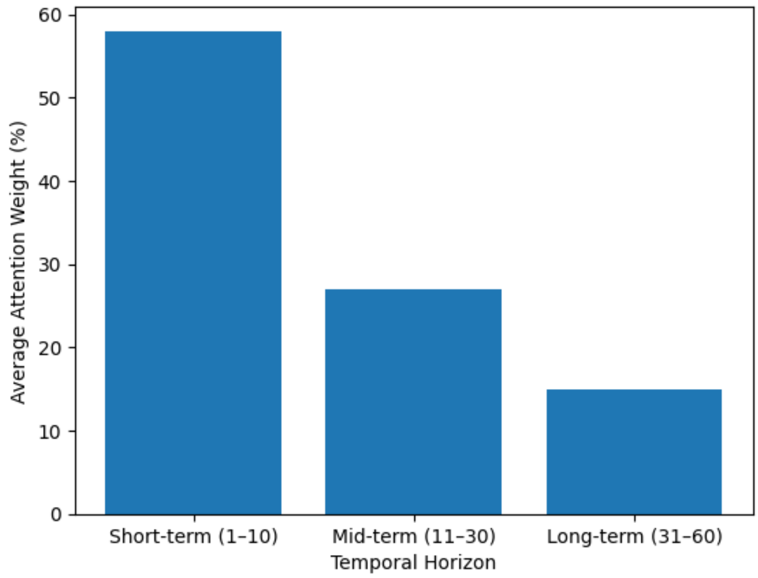

Figure 5. Distribution of Attention Weights Across Different Time Horizons

This figure (Figure 5) displays the distribution of the mean attention weights of the memory-instance-based gated Transformer model across three horizons—short-term, medium-term, and long-term—and reflects how temporal importance is allocated in the return forecasting process. The largest share of attention weight, approximately 58%, is assigned to the short-term horizon, indicating that the model is primarily sensitive to recent market changes and treats information close to the prediction time as the main driver of decision-making. This behavior is consistent with the nature of financial markets, where rapid responses to new information are critical, and it demonstrates the model's ability to extract short-term patterns and current market fluctuations. Nonetheless, the substantial weight assigned to the medium-term horizon (about 27%) indicates that the model does not rely solely on the most recent observations; rather, it uses more stable patterns from the preceding weeks to adjust its forecasts. The allocation of approximately 15% of attention weights to the long-term horizon directly reflects the role of instance-based memory in the MIGT structure. The presence of this meaningful share indicates that the model can incorporate more distant historical information—such as trends formed in prior months or the lingering effects of previous major market events—into the forecasting process. This capability is particularly important under regime-change conditions, because relying exclusively on short-term data can lead to unstable and overly reactive decisions. Therefore, the attention-weight distribution suggests that MIGT, by balancing multiple time horizons, maintains the flexibility needed to respond to rapid market developments while also leveraging instance-based memory to ensure stability and coherence of forecasts over longer horizons.

The evaluation was performed on out-of-sample test data using RMSE and MAE, in order to capture both average error magnitude and model sensitivity to large errors. The comparative results showed that MIGT achieved the lowest error values on both metrics: MAE = 0.0068 and RMSE = 0.0099. By contrast, the Long Short-Term Memory (LSTM) model yielded MAE = 0.0081 and RMSE = 0.0118, the Gated Recurrent Unit (GRU) model produced MAE = 0.0077 and RMSE = 0.0113, and the classic Transformer achieved MAE = 0.0072 and RMSE = 0.0106. Based on these results, MIGT improved upon the best competing model in terms of MAE (the classic Transformer) by 0.0004 and in terms of RMSE by 0.0007, which constitutes a meaningful reduction given the scale of daily returns. For a more precise interpretation, percentage improvements were also computed to enhance comparability. Relative to LSTM, the reduction in MAE from 0.0081 to 0.0068 corresponds to approximately a 16.05% improvement, and the reduction in RMSE from 0.0118 to 0.0099 corresponds to about a 16.10% improvement. Relative to GRU, the MAE reduction from 0.0077 to 0.0068 corresponds to 11.69%, and the RMSE reduction from 0.0113 to 0.0099 corresponds to 12.39%. Relative to the classic Transformer, the MAE reduction from 0.0072 to 0.0068 is estimated at 5.56%, and the RMSE reduction from 0.0106 to 0.0099 is estimated at 6.60%. A key methodological point is that the performance gap between MIGT and recurrent models (LSTM and GRU) is larger than the gap between MIGT and the classic Transformer. This pattern suggests that a major portion of MIGT's advantage derives from attention-based modeling of temporal dependencies, while instance-based memory and the gating mechanism provide an additional incremental improvement in accuracy.

Model behavior under severe volatility can also be inferred by comparing the RMSE–MAE gap. For MIGT, the ratio RMSE/MAE = 1.46, whereas the same ratio is 1.46 for LSTM, 1.47 for GRU, and 1.47 for the classic Transformer. The similarity of these ratios indicates that all models are affected by price jumps; however, the absolute error level of MIGT remains lower under these conditions. In other words, the proposed model not only reduces average error, but also partially controls the magnitude of large errors. This outcome is aligned with the analytical findings regarding attention and instance-based memory: attention highlights key time points, while instance-based memory stabilizes the influence of important past events in forecasts and prevents unstable reactions to noise. From an applied portfolio management perspective, even small reductions in forecast error—on the order of a few thousandths—can have a substantial effect on multi-objective optimization in the allocation stage, because the optimization algorithm makes decisions based on predicted-return signals, and any reduction in error reduces bias in expected-return estimates and improves weighting decisions.

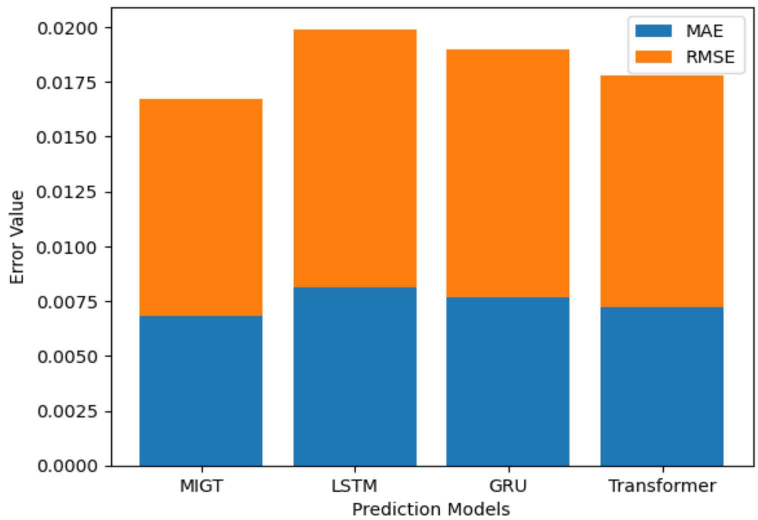

Figure 6. Comparison of Forecasting Errors Across Different Models

Figure 6 presents the performance of four forecasting models—MIGT, LSTM, GRU, and the classic Transformer—using the MAE and RMSE error metrics. As shown in the figure, MIGT achieves the lowest error values on both metrics, such that the total height of its corresponding bar is meaningfully lower than those of the other models. This indicates that MIGT not only has a lower average forecasting error, but also demonstrates more stable performance when confronted with large errors. In contrast, LSTM records the highest error values, followed by GRU and the classic Transformer. The simultaneous reduction of MAE and RMSE in MIGT suggests that, beyond improving average accuracy, the model more effectively controls the effects of severe volatility and extreme values—an essential feature for financial data characterized by heavy-tailed distributions and sudden shocks.

After data cleaning and preprocessing, the input data became statistically reliable for modeling, and the behavior of the time series was meaningfully linked to the quality of the final results. In the descriptive statistics stage, the mean daily returns of the selected assets were close to zero, and the dispersion of returns indicated that the market exhibits considerable volatility. The presence of pronounced skewness and kurtosis in a large portion of the assets showed that return distributions deviate from normality, which justifies the necessity of employing robust and multi-criteria risk measures. In addition, the initial correlation analysis revealed that some assets exhibit high co-movement; if diversification constraints and maximum weight limits are not imposed, there would be a risk of unintended concentration and increased idiosyncratic risk. These findings provided the scientific basis for defining weight and liquidity constraints, as well as for selecting downside and tail-risk measures in subsequent stages. In the experimental settings section, the results showed that the choice of operational parameters plays a direct role in the stability and reproducibility of outcomes. Implementing the forecasting and optimization models in the Python environment and running repeated experiments moved the evaluation beyond a single-run setting and increased the reliability of the results. Sequentially defining training, validation, and test windows prevented information leakage and ensured consistency with time-series logic. Moreover, aligning trading dates and managing missing data enabled the construction of a return matrix with minimal series breaks, leading to more stable estimation of risk measures (particularly the covariance matrix and tail-risk metrics). Finally, incorporating transaction costs of approximately 0.01 in return calculations ensured that the results were operationally realistic and that optimization was not conducted solely on gross returns.

The findings from the return forecasting phase showed that the memory-instance-based gated Transformer model achieved high accuracy in estimating asset returns and exhibited controlled sensitivity to market noise and instability. In the out-of-sample test period, the mean absolute error was 0.0068 and the root mean squared error was 0.0099, indicating satisfactory accuracy at the scale of daily returns. An important point was that the gap between MAE and RMSE suggested that large errors in the model output were constrained; thus, the model not only reduced average error but also partially mitigated the magnitude of extreme errors. This issue is particularly important for the optimization phase, because the optimization algorithm makes decisions based on predicted returns, and reducing large errors prevents the transmission of misleading signals to the weight-allocation stage. Comparison of the forecasting model with benchmark models showed that the superiority of the proposed model in error metrics was both significant and consistent. Relative to the Long Short-Term Memory model, MAE decreased from 0.0081 to 0.0068 and RMSE decreased from 0.0118 to 0.0099, representing an approximate 16% improvement. Compared with the Gated Recurrent Unit model, error reductions of about 12% were observed, and relative to the classic Transformer, although the gap was smaller, the error reduction remained observable and meaningful. This pattern indicated that the advantage of the proposed model operates on two levels: first, the use of the attention mechanism to learn temporal dependencies, and second, the addition of instance-based memory and gating to preserve key information and reduce the impact of noise.

An analysis of the model's performance in extracting temporal patterns showed that the attention mechanism and instance-based memory functioned in a complementary manner, each contributing differently to error reduction. The findings indicated that when the model concentrates more strongly on key temporal points, it is better able to incorporate the effects of important events into future forecasts, thereby reducing errors arising from regime changes. On the other hand, instance-based memory enabled the storage of historical event information with persistent effects, allowing the model to retain long-term dependencies when encountering similar periods. This feature is particularly important in markets characterized by recurrent shocks, because relying solely on short-term relationships can lead to unstable forecasts.

Discussion and Conclusion

The findings of this study demonstrate that the proposed Memory-Instance-based Gated Transformer (MIGT) framework delivers a statistically and economically meaningful improvement in asset return forecasting and portfolio management performance. The empirical results showed lower MAE and RMSE values for MIGT relative to benchmark models, along with improved directional accuracy and greater stability under volatile market conditions. These outcomes indicate that integrating gated attention mechanisms with instance-based memory enables the model to better capture the nonlinear, regime-dependent dynamics that characterize financial markets. This finding is consistent with prior evidence suggesting that traditional linear and purely recurrent models struggle to adapt to abrupt structural changes and heavy-tailed return distributions (1, 7). By contrast, attention-based architectures allow the model to focus selectively on informative historical segments, thereby improving robustness in the presence of noise and extreme observations.

From a theoretical perspective, the superior performance of MIGT can be explained by its alignment with both behavioral finance insights and modern machine-learning theory. Behavioral studies emphasize that markets are driven not only by fundamentals but also by sentiment, overreaction, and herding, which often manifest as clustered volatility and asymmetric return responses (2, 3). The attention mechanism in MIGT appears to operationalize this insight by dynamically reallocating importance across time, intensifying focus on key periods during turbulent regimes. This behavior mirrors empirical observations in emerging markets, including the Iranian capital market, where investor sentiment and macroeconomic shocks generate persistent effects rather than isolated price movements (4, 14). The instance-based memory component further reinforces this process by retaining salient historical patterns and reactivating them when similar conditions arise, thus preventing the loss of long-term dependencies.

The comparative analysis with benchmark models highlights an important methodological contribution of this study. While recurrent models such as LSTM and GRU are capable of modeling sequential dependencies, their reliance on fixed memory structures limits their ability to discriminate between noise and structurally important information, particularly over long horizons. The observed performance gap between MIGT and recurrent architectures supports earlier findings that attention-based models are better suited for financial time series characterized by non-stationarity and regime switching (10, 11). Moreover, although the classic Transformer also benefits from self-attention, the additional gains achieved by MIGT suggest that gating and memory augmentation play a critical role in stabilizing learning and reducing sensitivity to extreme events. This incremental improvement aligns with recent studies emphasizing the value of hybrid and robust learning architectures in financial forecasting tasks (5, 6).

The results also have important implications for portfolio optimization. Forecast accuracy alone is not sufficient; the stability and reliability of predictive signals directly affect downstream allocation decisions. The reduction in extreme forecasting errors observed for MIGT is particularly relevant in multi-objective optimization settings, where misestimated expected returns can lead to excessive concentration and heightened idiosyncratic risk. The findings of this study support the argument that improved predictive robustness enhances the effectiveness of portfolio constraints related to diversification, liquidity, and risk control (15, 16). In this sense, MIGT contributes not only to predictive performance but also to the economic feasibility and resilience of optimized portfolios.

Another notable aspect of the findings concerns interpretability. Although deep learning models are often criticized for their opacity, the attention structure within MIGT provides partial transparency by revealing which temporal segments exert the greatest influence on predictions. This characteristic addresses growing concerns about explainability in algorithmic finance and aligns with calls for more interpretable AI systems in regulated financial environments (12, 17). The observed regime-dependent redistribution of attention weights suggests that the model's decisions are not arbitrary but reflect economically meaningful shifts in market structure, such as transitions from calm to turbulent conditions.

Taken together, the empirical evidence indicates that MIGT offers a coherent and effective framework for portfolio management in complex financial environments. By combining attention-based learning, gated information flow, and instance-based memory, the model successfully integrates short-term responsiveness with long-term contextual awareness. This dual capability is particularly valuable in emerging and volatile markets, where historical shocks often reoccur and behavioral effects amplify market dynamics. The consistency of these findings with prior research on machine learning-based portfolio management reinforces the validity of the proposed approach and highlights its contribution to the evolving literature on intelligent financial decision-making (8, 9, 13).

Despite these contributions, several limitations of the present study should be acknowledged. The empirical analysis was conducted within a specific market context, which may limit the generalizability of the results to other financial systems with different regulatory structures, liquidity conditions, or investor compositions. In addition, although transaction costs were incorporated, other real-world frictions such as market impact, short-selling constraints, and tax considerations were not explicitly modeled. Finally, while attention weights provide partial interpretability, the internal representations of the memory module remain complex and warrant deeper analytical exploration.

Future research can extend this work in several directions. First, applying the MIGT framework to multiple international markets and cross-asset portfolios would provide a more comprehensive assessment of its generalizability. Second, integrating macroeconomic indicators, alternative data sources, and sentiment measures could further enhance the model's ability to anticipate regime shifts. Third, combining MIGT with advanced optimization techniques, such as risk-parity or sustainability-oriented objectives, may yield richer portfolio management insights under diverse investment mandates.

From a practical perspective, the findings of this study suggest that portfolio managers and financial institutions can benefit from adopting hybrid, attention-based learning systems that balance responsiveness and stability. Implementing such models within decision-support systems may improve return estimation, reduce exposure to extreme forecasting errors, and enhance the robustness of asset allocation strategies. Moreover, the selective and interpretable nature of attention mechanisms can support more transparent communication with stakeholders and regulators, thereby facilitating the responsible deployment of artificial intelligence in portfolio management practice.

Acknowledgments

We would like to express our appreciation and gratitude to all those who helped us carrying out this study.

Authors' Contributions

All authors equally contributed to this study.

Declaration of Interest

The authors of this article declared no conflict of interest.

Ethical Considerations

All ethical principles were adheried in conducting and writing this article.

Transparency of Data

In accordance with the principles of transparency and open research, we declare that all data and materials used in this study are available upon request.

Funding

This research was carried out independently with personal funding and without the financial support of any governmental or private institution or organization.

References

1. Antony A. Behavioral Finance and Portfolio Management: Review of Theory and Literature. Journal of Public Affairs. 2019;20(2). doi: 10.1002/pa.1996.
2. Hadbaa H. Behavioral Biases Influencing the Decision Making of Portfolio Managers of Capital Securities and Traders in Morocco. Financial Markets Institutions and Risks. 2019;3(1):92-105. doi: 10.21272/fmir.3(1).92-105.2019.
3. Bahramian M. Behavioral Portfolio Management: Tehran Stock Exchange Brokers Administration; 2022.
4. Montazeralhaj N, Rezaei Shouraki M, editors. Investigating the impact of investor sentiment on investment returns and financial performance in companies listed on the Tehran Stock Exchange2023.
5. Ayari Salah GH. A meta-analysis of supervised and unsupervised machine learning algorithms and their application to active portfolio management. Expert Systems with Applications. 2025;271:126611. doi: 10.1016/j.eswa.2025.126611.
6. Liu W, Suzuki Y, Du S. Forecasting the Stock Price of Listed Innovative SMEs Using Machine Learning Methods Based on Bayesian optimization: Evidence from China. Computational Economics. 2024;63(5):2035-68. doi: 10.1007/s10614-023-10393-4.
7. Han Y, Liu Y, Zhou G, Zhu Y. Technical analysis in the stock market: A review. Handbook of Investment Analysis, Portfolio Management, and Financial Derivatives. 1-42024. p. 1893-928.
8. Johnson L, Moore K. Reinforcement Learning for Financial Portfolio Management. IEEE Transactions on Neural Networks and Learning Systems. 2019;30(5):1306-17ER -.
9. Liang Z, Chen H, Zhu J, Jiang K, Li Y. Adversarial deep reinforcement learning in portfolio management2018.
10. Chen S, Huang Y, Ge L. An early warning system for financial crises: A temporal convolutional network approach. Technological and Economic Development of Economy. 2024;30(3):688-711. doi: 10.3846/tede.2024.20555.

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

12. Burkart N, Huber MF. A survey on the explainability of supervised machine learning. Journal of Artificial Intelligence Research. 2021;70:245-317. doi: 10.1613/jair.1.12228.

13. Wang N, Guo Z, Shang D, Li K. 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. 2024;200:123178. doi: 10.1016/j.techfore.2023.123178.

14. Rouhi Sara M, Taheri Nia M, Zolqi H, Sarlak A. Proposing a predictive model for financial crises in the Iranian capital market using hybrid algorithms. [Journal not specified]. 2023.

15. Drenovak M, Ranković V, Urošević B, Jelic R. Bond Portfolio Management Under Solvency II Regulation. European Journal of Finance. 2020;27(9):857-79. doi: 10.1080/1351847x.2020.1850499.

16. Silvius G, Marnewick C. Interlinking sustainability in organizational strategy, project portfolio management and project management a conceptual framework. Procedia Computer Science. 2022;196:938-47. doi: 10.1016/j.procs.2021.12.095.

17. Langley P, Leyshon A. FinTech platform regulation: regulating with/against platforms in the UK and China. Cambridge Journal of Regions, Economy and Society. 2023;16(2):257-68. doi: 10.1093/cjres/rsad005.