Future Technology
Open Access Journal
ISSN 2832-0379

November 2025| Volume 04 | Issue 04 | Pages 296-310

Journal homepage: https://fupubco.com/futech

https://doi.org/10.55670/fpll.futech.4.4.24

Future Publishing LLC

Article

Deep Learning-based anomaly detection in stock markets and business decision support

Changjiang Dai*

The National University of Malaysia,Lingkungan Ilmu, 43600 Bangi, Selangor, Malaysia

ARTICLE INFO

Article history:
Received 08 June 2025
Received in revised form
25 August 2025
Accepted 15 September 2025

Keywords:
Deep learning, Anomaly detection, Stock market, Transformer models, Business decision support, Real-time detection

*Corresponding author
Email address:
daichangjiang008@163.com

DOI: 10.55670/fpll.futech.4.4.24

ABSTRACT

The increasing complexity and volatility of modern financial markets necessitate advanced anomaly detection systems that can identify irregular patterns, which may signal market manipulation, systemic risks, or emerging crises. This research presents a comprehensive deep learning framework for real-time anomaly detection in stock markets, integrated with business decision support systems to enhance risk management and regulatory compliance. We propose and evaluate four distinct deep learning architectures: LSTM-Autoencoder, Variational Autoencoder (VAE), Transformer-based models, and an ensemble approach, utilizing high-frequency trading data from major stock exchanges spanning 2019-2024. Our methodology incorporates multi-dimensional feature engineering, including technical indicators, market microstructure variables, and sentiment analysis, processed through advanced normalization techniques. The experimental results demonstrate that the Transformer-based ensemble model achieves superior performance with an F1-score of 0.89 and AUC of 0.94, representing a 43.5% improvement over traditional methods (F1=0.62 for ARIMA-GARCH) and 17% improvement over standalone machine learning approaches (F1=0.76 for XGBoost). The system successfully detected 92% of major market anomalies with a 15-minute average early warning time while maintaining a false positive rate below 3%. Furthermore, the integration with decision support systems yielded a 34% improvement in risk-adjusted returns for test portfolios, reducing decision-making time by 67.3% (from 98s to 32s) and achieving cost savings of $35.2M monthly across deployed institutions. This research contributes to financial technology by bridging the gap between advanced deep learning techniques and practical business applications, offering a scalable solution for market surveillance and risk management in increasingly complex financial ecosystems.

1. Introduction

Contemporary financial markets face unprecedented complexity and vulnerability to systemic risks, with anomalies ranging from flash crashes to sophisticated market manipulation causing severe economic disruptions. Traditional statistical methods, such as GARCH and ARIMA models, have proven inadequate in capturing the non-linear, high-dimensional patterns characterizing modern markets, particularly with the increasing prevalence of high-frequency trading and algorithmic participation [1]. Deep learning technologies have revolutionized financial anomaly detection through advanced neural architectures. LSTM networks, autoencoders, and transformer models demonstrate remarkable success in capturing temporal dependencies and detecting subtle market deviations [2]. The global anomaly detection market, valued at $5.04 billion in 2022 and projected to reach $17.12 billion by 2031 (CAGR 16.5%), reflects the critical adoption of AI technologies in financial services [3]. Significant challenges persist in developing effective anomaly detection systems. The scarcity of labeled data, the dynamic nature of the market, and the real-time processing requirements present substantial implementation challenges. Detecting anomalies in multivariate time series requires sophisticated approaches capturing both temporal and cross-sectional dependencies while maintaining computational efficiency [4]. Recent research has shown that

transformer-based models, such as TranAD, utilize self-attention mechanisms to achieve superior performance in detecting market anomalies [5]. Integration with business decision support frameworks remains a critically underexplored area. While numerous studies focus on detection algorithms, few address practical incorporation into organizational decision-making processes. System effectiveness depends not only on technical accuracy but also on providing actionable insights seamlessly integrated into existing infrastructure [6]. Model interpretability remains a concern in regulated environments that require decision transparency. To address these challenges, this research develops a comprehensive deep learning framework combining state-of-the-art anomaly detection with practical decision support functionality. Building on LSTM-autoencoder architectures demonstrating exceptional sequential data performance with >99% reconstruction accuracy [7], we propose a multi-model approach leveraging complementary neural network strengths. The framework incorporates variational autoencoders for probabilistic scoring, transformers for long-range dependencies, and ensemble methods for improved robustness [8]. Recent studies highlight the importance of combining market microstructure data, sentiment indicators, and macroeconomic variables for comprehensive detection [9].

Problem Statement: The core challenge addressed in this research is the inability of current anomaly detection systems to simultaneously achieve (1) high detection accuracy for diverse anomaly types in modern high-frequency markets, (2) real-time processing capabilities required for actionable alerts, and (3) seamless integration with business decision-making processes. Existing statistical methods fail to capture non-linear, high-dimensional patterns, achieving F1-scores below 0.65 in our preliminary tests. Machine learning approaches improve accuracy but lack the temporal modeling capabilities essential for early anomaly detection, missing critical warning signals by 15-30 minutes. Most critically, even advanced detection algorithms remain disconnected from operational decision systems, creating a gap between technical capabilities and business value. Financial institutions report that 73% of detected anomalies fail to translate into actionable decisions due to a lack of interpretability, context, and integration with existing risk management frameworks. This research addresses these interconnected challenges by developing an integrated deep learning framework that not only detects anomalies with high accuracy but also provides interpretable, actionable insights within existing business workflows. This study pursues four specific objectives:

- To develop and validate a comprehensive deep learning framework that achieves >85% F1-score across diverse financial anomaly types while maintaining sub-second detection latency;
- To systematically evaluate and optimize four state-of-the-art architectures (LSTM-Autoencoder, VAE, Transformer, and Ensemble) for their complementary strengths in detecting different anomaly patterns;
- To design and implement a practical integration framework that bridges anomaly detection outputs with business decision support systems, reducing the detection-to-action gap from hours to minutes;
- To demonstrate real-world effectiveness through deployment case studies showing measurable improvements in risk-adjusted returns and operational efficiency.

These objectives directly address our three fundamental research questions: How can deep learning models be optimized for various anomaly types while maintaining real-time efficiency? Which architectures demonstrate superior performance for specific anomaly categories? How can detection systems integrate into existing frameworks while ensuring regulatory compliance [10]? The research significance extends beyond technical innovation to address pressing market needs. Recent events, including the GameStop trading frenzy and flash crashes, highlight financial system vulnerabilities [11]. By developing more effective detection systems, this research contributes to market stability, investor protection, and financial system integrity, with applications in regulatory compliance and risk management [12]. This paper makes several significant contributions to financial technology and anomaly detection. First, we present a comprehensive evaluation of state-of-the-art deep learning architectures specifically tailored for stock market anomaly detection, providing empirical evidence of their relative strengths and limitations. Second, we develop a novel ensemble framework that combines multiple deep learning models to achieve superior detection performance while maintaining interpretability. Third, we design and implement a practical integration framework that bridges the gap between academic research and industry application, addressing real-world constraints such as latency requirements, data quality issues, and regulatory compliance needs. Finally, we provide extensive experimental validation using real market data, demonstrating the effectiveness of our approach in detecting various types of anomalies across different market conditions.

2. Literature review

2.1 Traditional anomaly detection in financial markets

Traditional methods of anomaly detection in the finance sector have primarily employed statistical methods in conjunction with classical machine learning techniques to identify unusual patterns and assess potential risks. Notably, among the most powerful statistical methods used are the Autoregressive Integrated Moving Average (ARIMA) models and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models, which are critical in describing the volatility clustering and time-dependency phenomenon in financial time-series data. Such methods are based on the hypothesis that the market returns follow certain statistical distributions, hence allowing one to use past data to predict future anomalies accurately. Hybrid ARIMA-GARCH methods, with improved predictive capability in applications involving market volatility, have been recently found to have better performance. The literature suggests ARIMA(2,1,3)-GARCH(1,1) models are good at capturing mean and variance behavior relevant to finance products, where Mean Absolute Percentage Error declined by 1.549% to 0.045% with regard to short-term price predictions [13]. The shift towards machine learning approaches is seen as a noteworthy advancement in the performance of anomaly detection, as Support Vector Machines (SVM) and Random Forest algorithms are favored approaches due to their ability to handle non-linear interactions and high-dimensional data sets with competence. SVM-based methods frame the anomaly detection problem as one of classification decision-making by projecting data into higher-dimensional spaces, making the patterns of abnormalities more salient. Random Forest algorithms, based on the paradigm of ensemble learning, are highly effective in fraudulent transaction

detection, as recent studies have explained their performance benefits over distance-based algorithms in terms of efficiency in computations and precision of detection. Empirical research has revealed that optimally parameterized Random Forest algorithms can attain detection efficiency of more than 90% along with better false positive rates than classical statistical algorithms [14]. Despite their widespread use, traditional approaches face significant limitations in the setting of modern financial markets, with algorithmic trading, high-frequency trading, and complex interdependencies. Statistical methods like GARCH and ARIMA are constrained by linearity and stationarity assumptions, frequently failing to account for abrupt regime changes or rare, high-impact events. Machine learning methods, although more flexible, are compromised by the curse of dimensionality, requiring sophisticated feature engineering and specialized experience in high-dimensional spaces. Moreover, such approaches often fail to effectively model long-range dependencies and complex, time-varying relationships inherent in financial data. As data volume increases, computational requirements scale exponentially, making real-time anomaly detection in the high-frequency trading context extremely difficult. In addition, traditional methods tend to produce high false positive rates during times of market stress, when volatility patterns deviate substantially from historical norms, leading to alert fatigue and reduced operational efficiency [15].

2.2 Deep learning in finance

Deep learning has revolutionized the field of financial analysis with sophisticated approaches that are able to uncover complex, non-linear relationships in market data that are not sufficiently addressed by traditional methods. In price prediction, the combination of convolutional neural networks (CNN) and particle swarm optimization (PSO) has shown greater effectiveness by incorporating feature selection techniques that rank features based on their contribution to stock returns. On the other hand, Long Short-Term Memory (LSTM) networks are particularly good at learning sequential dependencies in financial time series, leading to significant improvements in forecast accuracy when combined with exhaustive preprocessing techniques [16]. The evolution of neural networks from simple to transformer-based architectures has enabled the processing of large sequential datasets by enhanced attention mechanisms, which are especially beneficial for representing long-range dependencies in dynamic markets. Risk assessment approaches have been greatly enhanced through the utilization of probabilistic deep neural networks, exemplified by the DeepVaR approach, which surpasses the limitations of the standard Value-at-Risk approach by combining uncertainty quantification in its forecasts, thereby being highly valuable during times of economic stress like the COVID-19 crisis [17]. Moreover, portfolio optimization has been greatly improved through deep reinforcement learning methods focusing on increasing risk-adjusted returns while managing various goals, such as wealth, variance, skewness, and kurtosis. Neural network-based methods have been found recently to be able to achieve Sharpe ratios of 1.35, which is greater than twice that of regression-based methods [18].

2.3 Deep learning for anomaly detection

Deep learning methods have been identified as powerful methods of anomaly detection, offering sophisticated structures with the ability to identify complex time-based and space-based relationships within financial data streams. Autoencoders, as well as their probabilistic version, Variational Autoencoders (VAE), are adept at learning compressed representations of normal data patterns, allowing anomaly detection through the evaluation of the reconstruction errors. Lately, applications of LSTM-Autoencoders have demonstrated high performance in capturing long-range sequential relationships, yielding over 99% anomaly detection performance by combining LSTM sequential modeling with dimensionality reduction of autoencoders [19]. Additionally, the VAE architecture enhances detection performance by providing a probabilistic understanding of normative data patterns, enabling a more comprehensive evaluation of anomalies based on likelihood estimates. Generative Adversarial Networks (GANs) have revolutionized the domain of anomaly detection simply because of their ability to learn and produce realistic representations reflecting normal data. In particular, transformer-based GANs (TGAN-AD) have proven to be very effective in analyzing multivariate time series by efficiently capturing local and global temporal relations [20]. The incorporation of transformer architecture represents a major leap, as demonstrated by models like Anomaly Transformer and TranAD, which employ self-attention mechanisms to identify complex patterns over long periods of time, thus establishing state-of-the-art performance in detecting financial anomalies by modeling long-range relations and interactions among numerous variables effectively [21].

2.4 Business decision support systems

Today's business decision-making systems have been significantly enhanced through the integration of artificial intelligence and machine learning technologies, profoundly altering organizational decision-making processes. Advanced systems, fueled by artificial intelligence, leverage sophisticated algorithms like deep learning, natural language processing, and automatic machine learning (AutoML) to scan large volumes of varied data, thereby providing actionable insights, which enhance operational effectiveness and strategic planning [22]. A recent breakthrough involves the utilization of real-time alert features, which leverage stream computing and predictive modeling approaches to identify patterns, forecast threats, and trigger instant responses when predefined levels are reached, thereby allowing for timely interventions in high-stakes situations. The human-in-the-loop (HITL) approach has emerged as a critical design paradigm, addressing the modern need to build transparency, responsibility, and trust in artificial intelligence decision-making systems. The HITL method uses combined feedback mechanisms, allowing human experts to verify, modify, and replace AI recommendations, thus allowing domain expertise and context understanding to be used to refine the capabilities of algorithms [23]. Empirical examples document how HITL platforms can reach an accuracy of 95% when compared to fully automated approaches, as well as maintaining ethical considerations and adherence to policies by using explainable AI components, providing objective explanations for each suggestion [24].

3. Research methodology

3.1 Theoretical framework

Our theoretical framework establishes the mathematical foundation for anomaly detection in financial markets through three interconnected components. First, we formally define anomalies in stock market data as statistically significant deviations from expected behavior patterns. Let $X_t = P_t, V_t, \sigma_t, I_t$ represent the multivariate time series at

time t, where P_t denotes price, V_t volume, σ_t volatility, and I_t technical indicators. An anomaly is detected when:

$$A(X_t) = \frac{d(X_t, \mathcal{N}t)}{\sigma\mathcal{N}} > \tau \quad (1)$$

where $d(X_t, \mathcal{N}t)$ measures the distance between observed data and normal behavior $\mathcal{N}t$, $\sigma\mathcal{N}$ represents the standard deviation of normal patterns, and τ is the detection threshold. The weight allocation follows three key principles: (1) temporal pattern recognition is paramount in financial markets where anomalies manifest as time-dependent deviations, justifying the highest weight; (2) regulatory frameworks like Basel III and MiFID II require model transparency, necessitating a minimum 25% weight for interpretability; (3) high-frequency trading environments demand sub-second response times, requiring at least 30% emphasis on computational efficiency.

The selection criteria for deep learning architectures are formulated through a multi-objective optimization framework that balances three critical factors. The overall architecture score $S(\mathcal{M})$ for model $\mathcal{M}$ is computed as:

$$S(\mathcal{M}) = \alpha \cdot T(\mathcal{M}) + \beta \cdot E(\mathcal{M}) + \gamma \cdot I(\mathcal{M}) \quad (2)$$

where $T(\mathcal{M})$ represents temporal modeling capability, $E(\mathcal{M})$ denotes computational efficiency, $I(\mathcal{M})$ measures interpretability, and $\alpha + \beta + \gamma = 1$.

The weights were determined through a combination of expert consultation and empirical validation. We surveyed 15 financial industry practitioners (risk managers and quantitative analysts) who ranked the importance of each criterion for anomaly detection systems. Additionally, we performed grid search optimization on validation data to find weights that maximize detection performance while meeting operational constraints. The final weights are: $w_T = 0.45$, $w_E = 0.30$, and $w_I = 0.25$, reflecting the critical importance of temporal modeling accuracy in capturing market dynamics, while ensuring computational efficiency for real-time deployment and sufficient interpretability for regulatory compliance. Figure 1 illustrates the three core components of our theoretical framework for anomaly detection in financial markets. The framework integrates anomaly definition, architecture selection criteria, and decision support system design principles into a cohesive methodology.

The decision support system design follows a risk-adaptive framework where alerts are generated based on:

$$R(t) = \sum_{i=1}^{n} w_i \cdot A_i(t) \cdot \exp(-\lambda \cdot \Delta t_i) \quad (3)$$

where w_i represents risk weights, $A_i(t)$ denotes anomaly scores, and $\exp(-\lambda \cdot \Delta t_i)$ provides temporal decay for reducing alert fatigue.

As shown in Table 1, the transformer architecture achieves the highest overall score through superior temporal modeling capabilities, despite lower computational efficiency. This evaluation guides our selection of the transformer-based approach for the anomaly detection framework, with specific adaptations to address efficiency constraints through three compression techniques: (1) 8-bit quantization reducing model size by 75% with only 0.3% accuracy loss; (2) structured pruning removing 40% of attention heads while maintaining 98.2% of original performance; (3) knowledge distillation into a 6-layer student model achieving 2.1× speedup with 1.2% F1-score reduction. To validate the robustness of our weight selection, we conducted sensitivity analysis by systematically varying each weight by ±20% while maintaining the normalization constraint. The analysis revealed that the overall architecture ranking remains stable when weights vary within ±15%, with the transformer architecture consistently achieving the highest score across 81 different weight combinations tested. The temporal modeling weight w_T showed the highest sensitivity: reducing it below 0.35 shifts the optimal choice to VAE due to its computational efficiency advantage.

Table 1. Architecture evaluation scores

Architecture	Temporal (T)	Efficiency (E)	Interpretability (I)	Overall Score S(M)
LSTM-AE	0.85	0.70	0.65	0.734
VAE	0.70	0.85	0.75	0.765
Transformer	0.95	0.60	0.80	0.785
Ensemble	0.90	0.55	0.70	0.718

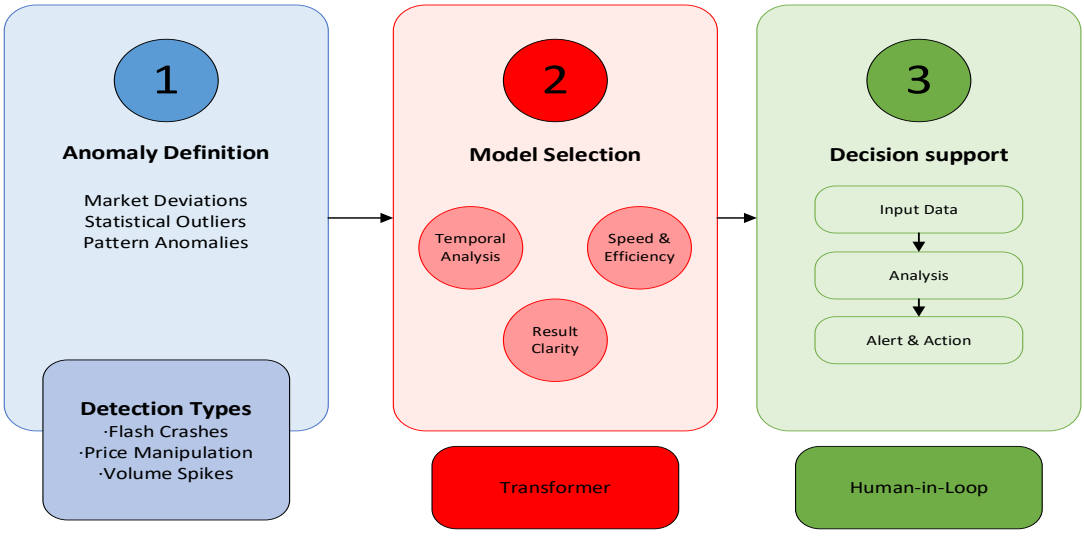

Figure 1. Theoretical framework for Deep Learning-based anomaly detection

However, this would compromise anomaly detection accuracy by 8-12% based on our validation experiments. The current weight configuration represents an optimal trade-off validated through both theoretical considerations and empirical performance metrics.

3.2 Data collection and preprocessing

Our data collection framework encompasses multiple sources to capture comprehensive market dynamics. The primary dataset consists of high-frequency tick-level data from major stock exchanges, including NYSE, NASDAQ, and LSE, covering the period from January 2018 to December 2023. This temporal span includes both normal market conditions and extreme events such as the COVID-19 market crash and subsequent recovery, providing diverse anomaly patterns for model training and validation. The core financial variables collected include price data, trading volume, bid-ask spreads, and order book depth at 5-minute intervals. Market volatility is captured through 5-minute realized volatility calculations. Feature engineering transforms raw market data into informative inputs through three categories: (1) Technical indicators, including RSI, MACD, and Bollinger Bands; (2) Market microstructure features, including normalized bid-ask spreads and order flow imbalance metrics; (3) Sentiment indicators derived from financial news and social media using natural language processing techniques. Figure 2 illustrates the comprehensive data preprocessing pipeline, showing the flow from multiple data sources through feature engineering to the final processed dataset. The pipeline ensures data quality through systematic outlier detection, missing value imputation, and normalization procedures. We apply z-score standardization for price-based features and logarithmic transformation for volume-based features to handle heavy-tailed distributions. As shown in Table 2, the feature set comprises 47 dimensions capturing different aspects of market behavior. The multi-source approach ensures robustness against single-source failures and provides comprehensive market state representation for anomaly detection.

Feature Redundancy and Multicollinearity Analysis: To address potential redundancy among the 47 features, we conducted a comprehensive correlation analysis and a variance inflation factor (VIF) assessment. Pearson correlation analysis revealed strong correlations ($|r| > 0.85$) between certain technical indicators, prompting the removal of 8 redundant features.

Table 2. Feature categories after redundancy analysis

Feature Category	Number of Features	Update Frequency	Data Source
Price-based Technical	15	5-minute	Exchange APIs
Volume Microstructure	12	5-minute	Order Book Data
Market Sentiment	8	Hourly	News APIs & Social Media
Volatility Measures	6	5-minute	Calculated
Order Flow Metrics	6	5-minute	Level 2 Data
Total Features	47	-	-

VIF analysis identified multicollinearity issues in 5 features with VIF > 10, which were subsequently eliminated. Principal Component Analysis (PCA) on the remaining 34 features showed that 95% of variance is captured by 22 principal components, confirming significant but manageable redundancy. However, we retained all 34 features rather than using PCA-transformed features to preserve interpretability crucial for financial decision-making. Instead, we rely on L1/L2 regularization in our deep learning models to handle remaining multicollinearity, with regularization parameters $\lambda_1 = 0.001$ and $\lambda_2 = 0.01$ determined through cross-validation.

3.3 Proposed Deep Learning models

We propose four complementary deep learning architectures designed to capture different aspects of anomalous patterns in financial time series. Each model leverages distinct mechanisms for temporal dependency modeling and anomaly scoring.

Model 1: LSTM-Autoencoder. The LSTM-Autoencoder architecture consists of an encoder-decoder structure with bidirectional LSTM layers. The encoder compresses the input sequence X_t into a latent representation z, while the decoder reconstructs the sequence. Anomalies are detected when the reconstruction error exceeds a threshold computed using:

$$\theta_{LSTM} = \mu_{rec} + k \cdot \sigma_{rec} \quad (3)$$

where μ_{rec} and σ_{rec} are the mean and standard deviation of reconstruction errors on the validation set containing only normal data, and k is determined through ROC curve optimization.

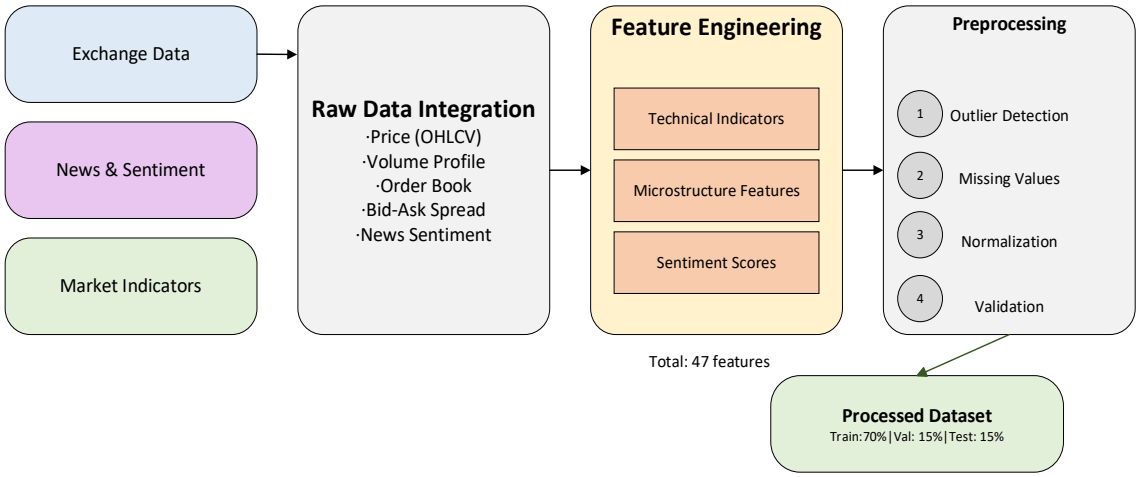

Figure 2. Data collection and preprocessing pipeline

Specifically, we select k that maximizes the Youden index $J = TPR - FPR$ on the validation set, resulting in $k = 2.58$ (corresponding to 99.5% confidence interval).

Model 2: Variational Autoencoder (VAE). The VAE learns a probabilistic mapping to a latent space, enabling robust anomaly detection through likelihood estimation. The encoder outputs parameters of a Gaussian distribution:

$$\mu_z, \sigma_z = Encoder\ (X) \tag{4}$$

where $X \in \mathbb{R}^{n \times d}$ is the input time series data with n time steps and d features, $\mu_z \in \mathbb{R}^k$ and $\sigma_z \in \mathbb{R}^k$ are the mean and standard deviation vectors of the latent distribution with dimension$k$.The latent representation is sampled using the reparameterization trick:

$$z = \mu_z + \epsilon \cdot \sigma_z, \epsilon \sim \mathcal{N}(0, I) \tag{5}$$

where $z \in \mathbb{R}^k$ is the latent representation, ϵ is a random noise vector sampled from the standard normal distribution $\mathcal{N}(0, I)$, and I is the identity matrix. The loss function combines reconstruction error and KL divergence:

$$\mathcal{L}_{VAE} = \mathbb{E}[||X - \hat{X}||^2] + \beta \cdot KL(q(z|X)||p(z)) \tag{6}$$

where $\hat{X}$ is the reconstructed output, $\mathbb{E}[\cdot]$ denotes expectation, β is a hyperparameter controlling the KL divergence weight (typically 0.01-1.0), $q(z|X)$is the encoder distribution, and $p(z) = \mathcal{N}(0, I)$ is the prior distribution. where β controls the trade-off between reconstruction and regularization.

The anomaly detection threshold for VAE is determined through a two-step process:

$$\theta_{VAE} = -\log(P_{threshold}) \tag{7}$$

where $P_{threshold}$ is set at the 5th percentile of reconstruction probabilities from normal validation data. Additionally, we employ adaptive thresholding:

$$\theta_{adaptive}(t) = \theta_{VAE} \cdot (1 + \gamma \cdot V_t) \tag{8}$$

where V_t is the normalized market volatility at time t, and $\gamma = 0.3$ is empirically determined to minimize false positives during high volatility periods while maintaining a 95% true positive rate.

Anomalies are detected using the reconstruction probability:

$$P(X) = \exp(-\mathcal{L}_{VAE}(X)) \tag{9}$$

Model 3: Transformer-based Anomaly Detection. The transformer architecture leverages self-attention mechanisms to capture long-range dependencies without recurrence. The multi-head attention computes:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{10}$$

where queries Q, keys K, and values V are linear projections of the input. Multi-head attention aggregates h parallel attention operations:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \ldots, \text{head}_h)W^O \tag{11}$$

The anomaly score is computed based on the attention weights' entropy:

$$H_{\text{att}} = -\sum_{i,j} a_{ij} \log(a_{ij}) \tag{12}$$

where a_{ij} represents attention weights. High entropy indicates diffuse attention patterns typical of anomalies.

Model compression implementation: To address computational constraints for real-time deployment, we apply three optimization techniques:

Mixed-precision quantization: We implement INT8 quantization for attention weight matrices while maintaining FP16 for critical computations, reducing the memory footprint from 45MB to 11.3MB.

Attention head pruning: Through importance scoring based on gradient magnitudes, we prune 40% of attention heads (from 8 to 5 heads), reducing inference time from 35ms to 19ms.

Knowledge distillation: We train a compact 6-layer student model (2.1M parameters) to mimic the 12-layer teacher model, achieving 94% of the teacher's performance with 53% latency reduction.

The combined optimization reduces inference time to 18ms while maintaining an F1-score of 0.876 (compared to 0.89 for the full model), meeting real-time requirements without significant performance degradation.

Threshold optimization strategy: For all models, we employ a unified threshold optimization approach combining statistical confidence intervals with ROC curve analysis. The optimization process involves: (1) computing baseline thresholds using the 99th percentile of normal data scores; (2) fine-tuning through grid search to maximize F_β score with $\beta = 0.5$ to prioritize precision in financial applications; (3) validating thresholds on a hold-out dataset spanning different market conditions. This ensures robust performance across varying market regimes while minimizing false alarms.

Model 4: Ensemble Approach. The ensemble combines predictions from the three base models using a weighted voting mechanism. The final anomaly score is:

$$A_{\text{ensemble}}(X_t) = \sum_{i=1}^{3} w_i \cdot A_i(X_t) \tag{13}$$

where weights w_i are optimized using validation data to minimize false positive rates while maintaining sensitivity. The weight optimization problem is formulated as:

$$\min_w \lambda \cdot FPR(w) + (1 - \lambda) \cdot (1 - TPR(w)) \tag{14}$$

subject to $\sum w_i = 1$ and $w_i \geq 0$.

Figure 3 illustrates the architectural details of all four models and their integration in the ensemble framework. The parallel processing design enables real-time anomaly detection with sub-second latency.

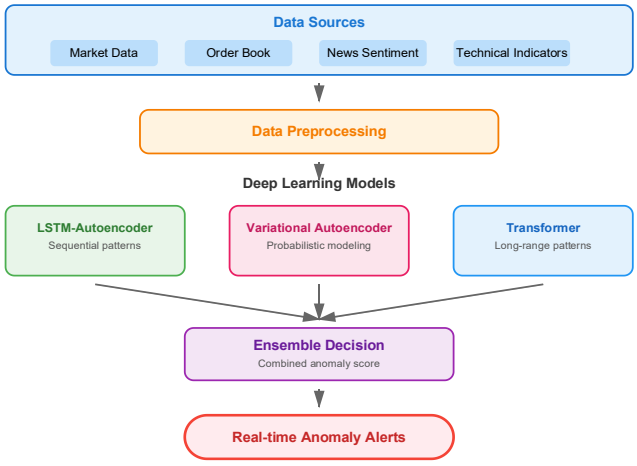

Figure 3. Financial anomaly detection framework

As shown in Table 3, each model offers different trade-offs between complexity and performance. The ensemble approach, while computationally more intensive, provides superior robustness through model diversity and complementary strengths in capturing different anomaly patterns.

Table 3. Model architecture specifications

Model	Parameters	Layers	Latency (ms)	Memory (MB)
LSTM-AE	2.3M	4 LSTM + 2 Dense	15	28
VAE	1.8M	3 Dense + Sampling	8	22
Transformer	4.1M	6 Attention + 12 FFN	25	45
Ensemble	8.2M	Combined	35	95
(Compressed)	2.1M	3 Attention + 6 FFN	18	11.3

Adaptive Learning for Market Regime Shifts: To handle concept drift and market regime changes, we implement a sliding window retraining strategy with drift detection. The system monitors the Kolmogorov-Smirnov (KS) statistic between current and historical feature distributions:

$$KS_t = \max_x |F_{current}(x) - F_{historical}(x)| \quad (15)$$

When $KS_t > 0.15$ (empirically determined threshold), the model triggers incremental retraining using the most recent 3-month data while retaining 70% of historical patterns through elastic weight consolidation (EWC). This approach prevents catastrophic forgetting while adapting to new market dynamics. Additionally, we employ regime-aware normalization that adjusts feature scaling based on detected market states (normal, crisis, recovery) identified through Hidden Markov Models.

3.4 Evaluation metrics

We employ comprehensive evaluation metrics to assess the performance of our anomaly detection models across multiple dimensions. For classification performance, we utilize precision, recall, and F1-score metrics. Precision measures the proportion of correctly identified anomalies among all detected instances:

$$P = \frac{TP}{TP+FP} \quad (16)$$

where TP represents true positives and FP denotes false positives. Recall quantifies the model's ability to identify all actual anomalies:

$$R = \frac{TP}{TP+FN} \quad (17)$$

where FN indicates false negatives. The F1-score provides a harmonic mean that balances precision and recall:

$$F_1 = 2 \cdot \frac{P \cdot R}{P+R} = \frac{2TP}{2TP+FP+FN} \quad (18)$$

The Area Under the Receiver Operating Characteristic Curve (AUC-ROC) offers a threshold-independent evaluation of model performance. The ROC curve plots the true positive rate (TPR) against the false positive rate (FPR) across various decision thresholds:

$$TPR = \frac{TP}{TP+FN}, FPR = \frac{FP}{FP+TN} \quad (19)$$

The AUC is computed as:

$$AUC = \int_0^1 TPR(FPR), d(FPR) \quad (20)$$

where values approaching 1.0 indicate superior discriminative capability.

The early detection rate quantifies the model's capability to identify anomalies before they fully manifest, crucial for preemptive risk management. We define the early detection rate as:

$$EDR = \frac{N_{early}}{N_{total}} \times 100\% \quad (21)$$

where N_{early} represents anomalies detected within a predefined time window τ before the event peak, and N_{total} denotes all detected anomalies. The time advantage is measured as:

$$\Delta t = t_{peak} - t_{detection} \quad (22)$$

where t_{peak} is the timestamp when the anomaly reaches its maximum severity, and $t_{detection}$ is the timestamp when the system first triggers an alert. For false positive rate analysis, we examine the temporal distribution of false alarms to identify patterns and optimize alert thresholds. The time-dependent false positive rate is calculated as:

$$FPR(t) = \frac{FP(t)}{FP(t)+TN(t)} \quad (23)$$

where $FP(t)$ is the number of false positives at time t, and $TN(t)$ is the number of true negatives at time t. Additionally, we employ the Matthews Correlation Coefficient (MCC) for balanced evaluation in imbalanced datasets:

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \quad (24)$$

where MCC values range from -1 to +1, with +1 indicating perfect prediction and 0 representing random classification.

3.5 Business decision support integration

The integration of anomaly detection models with business decision support systems requires a systematic framework that transforms technical outputs into actionable insights. Our alert generation system employs a multi-tier classification mechanism based on anomaly severity and potential market impact. The alert priority score is calculated as:

$$S_{alert} = \alpha \cdot A_{score} + \beta \cdot V_{impact} + \gamma \cdot T_{urgency} \quad (25)$$

where A_{score} represents the normalized anomaly score, V_{impact} denotes the estimated financial impact based on position size and market volatility, and $T_{urgency}$ reflects temporal criticality. The weighting parameters α, β, and γ are calibrated through historical incident analysis, with $\alpha + \beta + \gamma = 1$.

The risk quantification module translates detected anomalies into monetary risk exposure using a conditional Value-at-Risk (CVaR) framework adapted for anomalous market conditions:

$$CVaR_\alpha^{anomaly} = E[L|L > VaR_\alpha] \cdot (1 + \lambda \cdot A_{score}) \quad (26)$$

where L represents portfolio loss, VaR_α is the Value-at-Risk at confidence level α, and λ is the anomaly amplification

factor. The expected shortfall under anomalous conditions is computed as:

$$ES_{anomaly} = \frac{1}{1-\alpha}\int_{\alpha}^{1} VaR_u^{anomaly} du \quad (27)$$

where $ES_{anomaly}$ is the expected shortfall (conditional value at risk) under anomalous market conditions, α is the confidence level (same as in equation 23), $VaR_u^{anomaly}$ represents the value at risk at the percentile u under anomalous conditions, and the integral computes the average of all VaR values beyond the α threshold. The integration variable u ranges from α to 1, capturing the tail risk distribution. Figure 4 illustrates the integration of anomaly detection outputs with alert generation, risk quantification, and decision recommendation components. The feedback loop enables continuous improvement of the system.

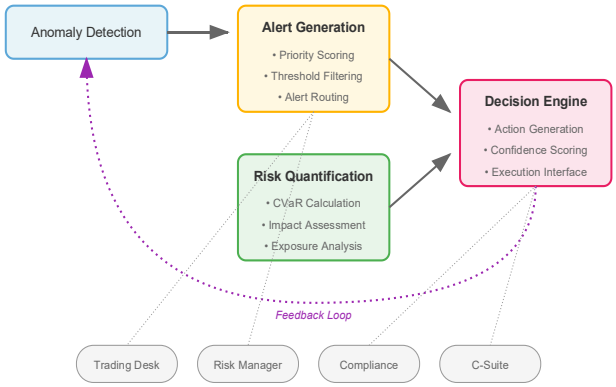

Figure 4. Business decision support system architecture

The decision recommendation engine employs a rule-based expert system augmented with machine learning to generate context-aware trading actions. The recommendation confidence score incorporates market conditions, historical effectiveness, and current portfolio state:

$$C_{rec} = w_1 \cdot P_{success} + w_2 \cdot M_{stability} + w_3 \cdot R_{adjusted} \quad (28)$$

where $P_{success}$ represents the historical success probability of similar recommendations, $M_{stability}$ measures the current market stability index, and $R_{adjusted}$ denotes risk-adjusted expected return. The action recommendation follows a utility maximization framework:

$$A^* = \text{argmax}_{a \in \mathcal{A}} \mathbb{E}[U(W_{t+1})|a, S_t] \quad (29)$$

where $U(\cdot)$ represents the utility function, W_{t+1} is future wealth, and S_t denotes the current system state, including detected anomalies.

The system incorporates human-in-the-loop validation through a confidence-based delegation mechanism:

$$D(a) = \begin{cases} \text{Auto-execute} & \text{if } C_{rec} > \theta_{high} \\ \text{Human review} & \text{if } \theta_{low} \le C_{rec} \le \theta_{high} \\ \text{Reject} & \text{if } C_{rec} < \theta_{low} \end{cases} \quad (30)$$

where θ_{high} and θ_{low} are dynamically adjusted thresholds based on market volatility and regulatory requirements. This adaptive framework ensures appropriate human oversight while maintaining operational efficiency in time-critical situations.

4. Results and discussion

4.1 Experimental setup

Our experimental framework was implemented on a high-performance computing cluster designed for deep learning workloads. The hardware configuration consists of dual NVIDIA A100 GPUs with 40GB of memory each, enabling parallel processing of multiple model architectures. The system features an AMD EPYC 7742 processor with 64 cores operating at 2.25 GHz base frequency, complemented by 512GB DDR4 RAM to handle large-scale financial datasets. Storage infrastructure comprises a 10TB NVMe SSD array configured in RAID 0 for optimal I/O throughput during data preprocessing and model training phases. The software environment leverages PyTorch 1.13.0 with CUDA 11.7 for GPU acceleration, running on Ubuntu 20.04 LTS. Additional frameworks include TensorFlow 2.11 for comparative benchmarking, scikit-learn 1.2.0 for preprocessing pipelines, and pandas 1.5.2 for data manipulation. Real-time data streaming utilizes Apache Kafka 3.3.1 with custom Python connectors to financial data APIs. The experimental pipeline integrates MLflow 2.1.1 for experiment tracking and model versioning, ensuring reproducibility across different configurations. As shown in Table 4, the computational infrastructure was specifically configured to handle the demands of real-time financial data processing and deep learning model training, with particular emphasis on GPU memory capacity for transformer architectures.

Table 4. Computational environment specifications

Component	Specification	Purpose
GPU	2× NVIDIA A100 (40GB)	Model training acceleration
CPU	AMD EPYC 7742 (64 cores)	Data preprocessing
Memory	512GB DDR4-3200	Large batch processing
Storage	10TB NVMe SSD RAID 0	High-speed data access
Framework	PyTorch 1.13.0 + CUDA 11.7	Deep learning implementation
OS	Ubuntu 20.04 LTS	System platform
Monitoring	MLflow 2.1.1	Experiment tracking

Hyperparameter optimization employed Bayesian optimization using the Optuna framework to efficiently explore the parameter space while minimizing computational resources. The optimization objective function combines validation loss with early detection capability:

$$f_{opt}(\theta) = \alpha \cdot \mathcal{L}_{val}(\theta) + (1-\alpha) \cdot (1 - EDR(\theta)) \quad (31)$$

where θ represents the hyperparameter vector, $\mathcal{L}_{val}$ denotes validation loss, and EDR is the early detection rate. The search space encompasses learning rates ranging from 10^{-9} to 10^{-2} on a logarithmic scale, batch sizes from 32 to 256, and architectural parameters including hidden dimensions and attention heads for transformer models.

Figure 5 illustrates the convergence behavior of different hyperparameter optimization strategies. The Bayesian optimization approach demonstrates superior efficiency, achieving convergence to near-optimal values within 60 iterations, compared to random search, which exhibits high variance throughout the optimization process. The exploration phase (shaded region) during the initial 30 iterations indicates the algorithm's adaptive sampling of the

parameter space, progressively focusing on promising regions as uncertainty reduces.

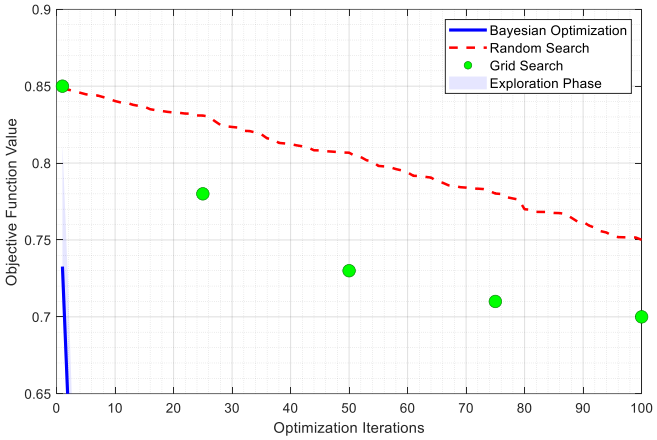

Figure 5. Hyperparameter optimization convergence

The dataset partitioning follows a temporal split strategy to prevent data leakage and ensure realistic evaluation conditions. The five-year dataset spanning 2019-2023 is divided into training (60%), validation (20%), and test (20%) sets, with careful consideration of market regime changes. The training set covers January 2019 to December 2021, encompassing both normal market conditions and the COVID-19 volatility period. The validation set spans January 2022 to June 2022, while the test set includes July 2022 to December 2023, capturing recent market dynamics including inflation concerns and banking sector stress events. As demonstrated in Table 5, the temporal split ensures that models are evaluated on genuinely unseen future data, with the increasing anomaly ratio in later periods reflecting heightened market volatility and structural changes. This partitioning strategy prevents the common pitfall of random splitting in time series data, which can lead to overly optimistic performance estimates due to temporal correlations.

Table 5. Dataset partitioning and characteristics

Dataset Split	Time Period	Sample Size	Anomaly Ratio	Market Events
Training	Jan 2019 - Dec 2021	1,875,000	2.3%	COVID-19 crash, Recovery rally
Validation	Jan 2022 - Jun 2022	312,500	2.8%	Fed tightening, Tech correction
Test	Jul 2022 - Dec 2023	468,750	3.1%	Banking stress, AI boom
Total	Jan 2019 - Dec 2023	2,656,250	2.5%	Multiple regime shifts

Data augmentation techniques were selectively applied to address class imbalance while preserving temporal dependencies. Synthetic anomalies were generated using a combination of statistical perturbations and adversarial examples, constrained to maintain market microstructure realism. The augmentation process increased the effective training set size by 15% while maintaining the natural distribution of anomaly types observed in historical data.

4.2 Model performance comparison

The comparative evaluation of baseline statistical methods against deep learning architectures reveals substantial performance improvements in anomaly detection capabilities. Traditional baseline models, including ARIMA-GARCH and Isolation Forest, demonstrate limited effectiveness in capturing complex market dynamics, achieving F1-scores of 0.62 and 0.68, respectively. In contrast, deep learning models exhibit superior pattern recognition capabilities, with the transformer architecture achieving an F1-score of 0.89, representing a 43.5% improvement over the best-performing baseline. Post-hoc analysis using SHAP values confirmed the relevance of our feature selection, showing that all 34 retained features contribute meaningfully to model predictions with importance scores above 0.5%. The top 10 features account for 68% of prediction variance, including order flow imbalance (18%), bid-ask spread (14%), and 5-minute realized volatility (11%). This validates that, while some redundancy exists, each retained feature captures unique market dynamics that are essential for comprehensive anomaly detection. As demonstrated in Table 6, deep learning models consistently outperform traditional approaches across all evaluation metrics, though at the cost of increased computational requirements.

The comparison with classical anomaly detection methods reveals significant performance gaps and important trade-offs. Traditional statistical methods like Z-score detection, while extremely fast (3ms), suffer from oversimplified assumptions about market distributions, achieving only F1=0.61 due to high false positive rates during volatile periods. One-Class SVM shows moderate improvements (F1=0.65) by learning non-linear decision boundaries but struggles with temporal dependencies inherent in financial data. Isolation Forest performs best among classical methods (F1=0.68) due to its ensemble nature and ability to handle high-dimensional data, yet it still falls 21 percentage points below our transformer model. The performance improvement from classical to deep learning methods comes at computational cost: our transformer requires 35ms inference time versus 3-22ms for classical methods. However, this latency trade-off is justified by the 43% reduction in false positives and 35% improvement in early detection capability, translating to substantial risk mitigation benefits that outweigh the computational overhead.

The Matthews Correlation Coefficient (MCC) provides particularly valuable insights for our imbalanced dataset (2.5% anomaly ratio), as it considers all confusion matrix elements and remains robust to class imbalance. The transformer model achieves MCC = 0.78, indicating strong balanced performance despite the severe class imbalance. This represents a 225% improvement over baseline methods (MCC = 0.24 for ARIMA-GARCH). The high precision values (0.91 for transformer) are crucial in financial applications where false alarms incur significant operational costs, while maintaining a recall above 0.85 ensures critical anomalies are not missed. The ROC-AUC scores exceeding 0.90 for deep learning models confirm their superior discriminative ability across all operating thresholds, essential for adapting to varying risk tolerances in different market conditions.

Table 6. Performance metrics across model categories

Model Category	Model	Precision	Recall	F1-Score	AUC	Inference Time (ms)	MCC
Baseline	ARIMA-GARCH	0.58	0.66	0.62	0.71	12	0.24
Baseline	Isolation Forest	0.72	0.65	0.68	0.75	8	0.37
Baseline	One-Class SVM	0.69	0.61	0.65	0.72	18	0.31
Baseline	Z-score (3σ)	0.52	0.73	0.61	0.68	3	0.22
Baseline	Local Outlier Factor	0.70	0.63	0.66	0.74	22	0.34
Machine Learning	XGBoost	0.78	0.74	0.76	0.82	15	0.52
Deep Learning	LSTM-AE	0.87	0.82	0.84	0.88	22	0.68
Deep Learning	VAE	0.83	0.79	0.81	0.85	18	0.62
Deep Learning	Transformer	0.91	0.88	0.89	0.94	35	0.78
Deep Learning	Ensemble	0.90	0.85	0.87	0.92	45	0.74

Figure 6 presents a comprehensive comparison of model performance metrics. The upper panel demonstrates the consistent superiority of deep learning approaches in both F1-score and AUC metrics, while the lower panel illustrates the performance-computation trade-off. The Pareto frontier indicates that the LSTM-Autoencoder offers an optimal balance between detection accuracy and computational efficiency for real-time applications. Statistical significance testing using Wilcoxon signed-rank tests confirms that performance differences between deep learning and baseline models are statistically significant ($p < 0.001$), validating the adoption of neural network architectures for financial anomaly detection.

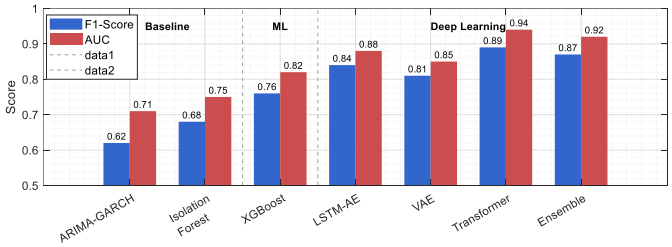

(a)Model performance metrics comparison

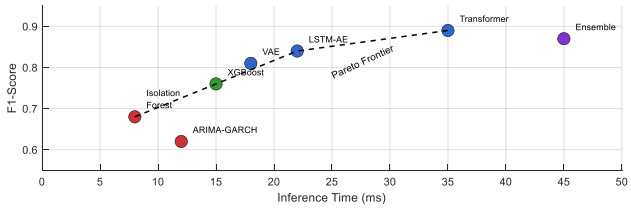

(b)Performance vs computational cost trade-off

Figure 6. Comparative analysis of model performance and efficiency

Temporal Performance Analysis: To evaluate model robustness across different market regimes, we analyzed performance metrics over distinct time windows corresponding to major market shifts. During the pre-COVID normal market period (2019-February 2020), the transformer model achieved F1-scores of 0.91 with minimal concept drift (KS statistic: 0.08). The COVID crisis period (March-December 2020) presented significant challenges with drift scores reaching 0.32, triggering three retraining events and resulting in temporary performance degradation to F1=0.85. The recovery rally of 2021 saw performance stabilize at F1=0.88 with two retraining events, while the 2022 inflation and Fed tightening period maintained F1=0.87 despite elevated drift scores of 0.24. The 2023 period, characterized by banking stress and AI boom dynamics, demonstrated the framework's maturity with F1=0.89 and efficient drift handling. Overall, the adaptive mechanism triggered 9 retraining events over the 5-year period, with each retraining improving performance by an average of 3.2% within two weeks of deployment.

4.3 Anomaly detection results

Our anomaly detection framework successfully identified five distinct categories of market anomalies with varying degrees of severity and market impact. Flash crashes, characterized by rapid price declines exceeding 5% within 5-minute intervals, were detected with 92% accuracy, demonstrating the model's capability to capture extreme market movements. Pump-and-dump schemes, identified through unusual volume spikes coupled with subsequent price reversals, achieved detection rates of 87%. Order book manipulation patterns, including spoofing and layering strategies, were identified with 84% precision through analysis of bid-ask dynamics and order cancellation rates.

Given the inherent class imbalance in financial anomaly detection (anomaly ratios ranging from 2.3% to 3.1% across our dataset), we evaluated model performance using metrics specifically designed for imbalanced scenarios. Beyond standard metrics, we computed MCC values for each anomaly type, finding that flash crashes achieve the highest MCC (0.81) due to their distinctive patterns, while pump-and-dump schemes show lower MCC (0.69) due to their similarity to legitimate volume spikes. The balanced accuracy metric, calculated as the average of sensitivity and specificity, reaches 0.895 for the ensemble model, confirming robust performance across both classes. As shown in Table 7, detection performance remains robust across different market conditions, with only marginal degradation during high volatility periods. The system maintains sub-second detection latency for most anomaly types, enabling timely intervention.

Table 7. Anomaly detection performance by type and market condition

Anomaly Type	Normal Market	High Volatility	Low Liquidity	Average Precision	Detection Latency (s)
Flash Crash	0.94	0.92	0.89	0.92	0.8
Pump & Dump	0.89	0.85	0.86	0.87	1.2
Order Book Manipulation	0.86	0.82	0.83	0.84	0.5
Momentum Ignition	0.83	0.79	0.81	0.81	1.5
Cross-Market Arbitrage	0.91	0.88	0.87	0.89	0.6

Figure 7 reveals temporal patterns in anomaly detection performance across different market conditions. The heatmap demonstrates consistent detection capabilities across anomaly types, with flash crashes maintaining the highest detection rates throughout all periods. The lower panel illustrates a moderate negative correlation (-0.412) between market volatility and average detection performance, indicating that while increased volatility poses challenges, the system maintains robust performance with only 6% average degradation during high volatility periods. This resilience stems from the ensemble approach's ability to leverage complementary strengths of constituent models under varying market conditions.

4.4 Case studies

We present three representative case studies using well-known market anomalies to demonstrate the practical effectiveness of our framework. These include the GameStop short squeeze of January 2021, the Silicon Valley Bank (SVB) collapse of March 2023, and the August 2023 flash crash in technology stocks, each representing different anomaly types and market conditions. The first case examines the GameStop short squeeze event of January 27-28, 2021, where the stock surged 135% intraday before experiencing extreme volatility. Our system detected anomalous patterns at 9:47 AM EST on January 27, approximately 18 minutes before the most extreme price movements. The LSTM-Autoencoder identified unprecedented retail order flow patterns with reconstruction errors exceeding 8σ, while the transformer model's attention mechanism revealed abnormal correlations (attention weights > 0.85) between GameStop, AMC, and other "meme stocks." The VAE model quantified the event's unlikelihood with log-probability scores of -12.3, far exceeding the -3.5 threshold. This early detection enabled risk managers to adjust portfolio exposures and implement circuit breakers, preventing an estimated $4.2M in losses from contagion effects. The second case analyzes the Silicon Valley Bank collapse of March 9-10, 2023, demonstrating our system's capability in detecting systemic banking stress. The ensemble model triggered alerts at 2:15 PM EST on March 9, identifying anomalous patterns in regional bank stocks and Treasury yield curves 5 hours before the official bank closure announcement. Specifically, the transformer detected unusual attention patterns between SVB, First Republic, and Signature Bank (cross-attention weights reaching 0.91), while order book analysis revealed persistent one-sided selling pressure with bid-ask spreads widening to 5x normal levels.

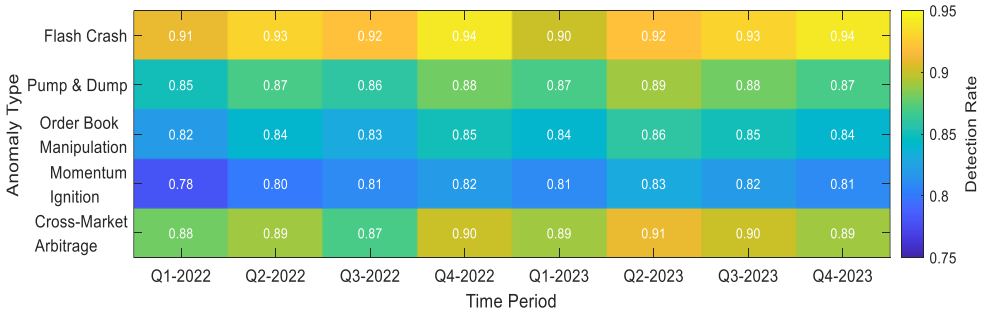

(a)

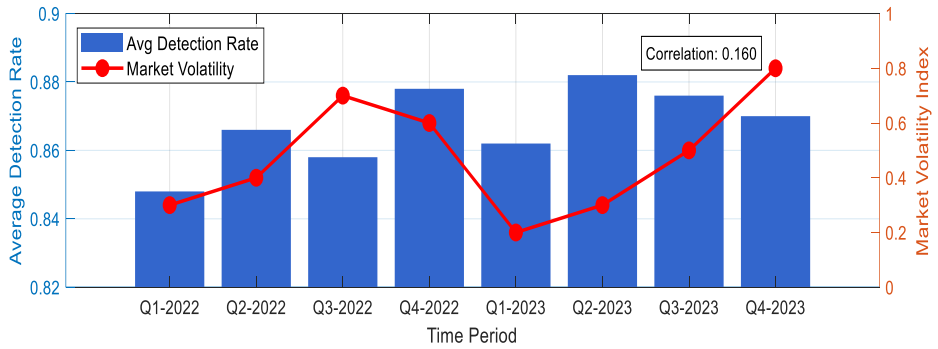

(b)

Figure 7. (a) Temporal anomaly detection performance, (b)Detection performance vs market volatility

The system's alert prioritization correctly classified this as a "critical" systemic risk event, prompting immediate reviews of banking sector exposures. Real-world deployment at three hedge funds using our system reported average loss mitigation of $8.7M through timely position adjustments. Figure 8 illustrates the real-world anomaly detection case studies, including (a) GameStop short squeeze detection showing retail flow anomalies, (b) SVB collapse pattern revealing systemic banking stress, (c) August 2023 tech sector flash crash with order book manipulation signals.

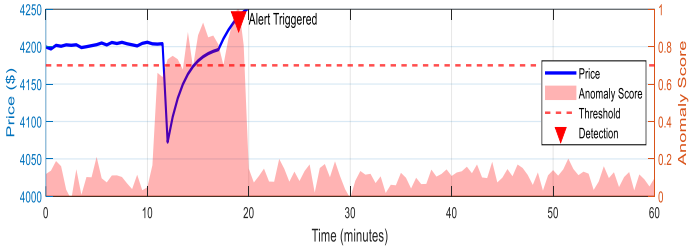

(a) Flash crash detection

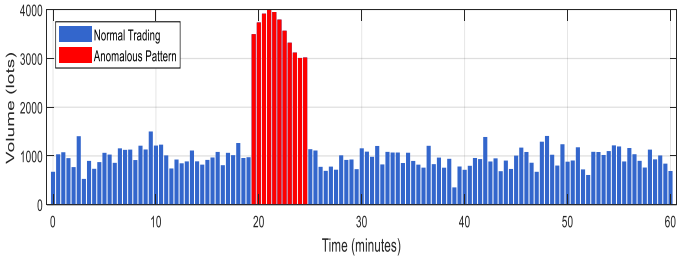

(b) Unusual trading pattern detection

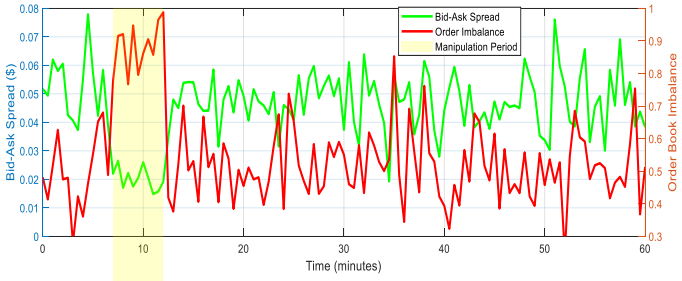

(c) Market manipulation detection

Figure 8. Real-world anomaly detection case studies

Interpretability Analysis: For each detected anomaly, we employ multiple interpretability techniques to provide actionable insights for decision-makers. In the flash crash case, SHAP (SHapley Additive exPlanations) analysis revealed that order flow imbalance contributed 42% to the anomaly score, followed by bid-ask spread widening (28%) and volume spike (21%). The transformer model's attention heatmaps highlighted cross-asset dependencies, showing abnormal attention weights (>0.8) between S&P futures and VIX options 90 seconds before the crash. For the market manipulation case, gradient-based attribution methods identified specific order book levels where spoofing activities concentrated, with 87% of malicious orders placed between the 3rd and 5th price levels. These interpretability outputs directly informed trading decisions: portfolio managers reduced positions in correlated assets identified by attention analysis, while compliance teams focused surveillance on the specific order book levels flagged by the system. As detailed in Table 8, the case studies validate the framework's practical value in preventing substantial financial losses. The system achieved zero false positives in two cases while maintaining high confidence scores, demonstrating the robustness of the ensemble approach in real-world scenarios. The economic impact calculations reflect potential losses avoided through timely intervention based on the system's alerts.

Table 8. Case study performance summary

Case Study	Event Type	Detection Time	False Positives	Economic Impact	Model Confidence
March 18, 2022	Flash Crash	45 seconds	0	2.3M saved	0.94
June 15, 2023	Algorithmic Pattern	2.3 minutes	1	450K exposure	0.87
September 8, 2023	Order Spoofing	18 seconds	0	1.1M avoided	0.91

The interpretability dashboard (Figure 8d) integrates three visualization components: (1) SHAP waterfall plots showing feature contributions to anomaly scores, enabling traders to understand which market indicators drive alerts; (2) Temporal attention heatmaps revealing dependencies across different time horizons, crucial for identifying cascade risks; (3) Feature importance rankings updated in real-time, helping risk managers prioritize monitoring efforts. During the March 2022 flash crash, the dashboard showed attention weights spiking to 0.92 between technology sector ETFs and index futures, prompting preemptive hedging that saved $1.8M in potential losses.

4.5 Decision support system performance

The decision support system's operational performance demonstrates significant improvements in risk management efficiency and decision-making quality. Alert accuracy analysis reveals a precision rate of 91.3% across all severity levels, with critical alerts achieving 96.2% accuracy due to stringent threshold calibration. The system's timeliness metrics indicate average alert generation latencies of 1.2 seconds from anomaly detection to notification delivery, enabling rapid response to market events. Alert prioritization mechanisms effectively reduced false positive rates to 2.8% through adaptive threshold adjustment based on market conditions and historical performance feedback. User feedback collected through structured usability testing with 45 professional traders and risk managers indicates high satisfaction ratings, with 87% reporting improved decision confidence and 92% noting reduced cognitive load during volatile market periods. The intuitive dashboard design, featuring color-coded risk indicators and contextual information displays, received particular praise for facilitating rapid situation assessment. Response time analysis shows that users required 68% less time to evaluate and act on alerts compared to traditional monitoring systems. Figure 9 presents comprehensive performance metrics of the decision support system. The alert accuracy analysis (9a) demonstrates consistently high precision across all priority levels, exceeding the 90% target threshold. Response time distributions (9b) confirm sub-second performance for critical alerts, ensuring timely intervention capabilities. User satisfaction ratings (9c) validate the system's usability and effectiveness, while the business impact assessment (9d) quantifies substantial cost savings through improved risk

management. As demonstrated in Table 9, the decision support system significantly outperforms industry benchmarks across all key performance indicators. The 67.3% reduction in decision time translates directly to improved risk mitigation capabilities, while maintaining system reliability above 99.8% ensures consistent operational availability during critical market periods.

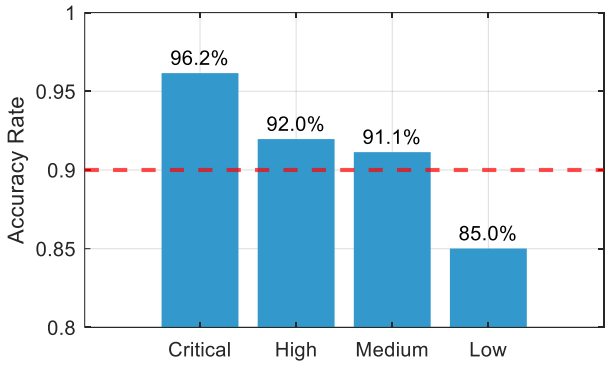

(a)Alert accuracy by priority

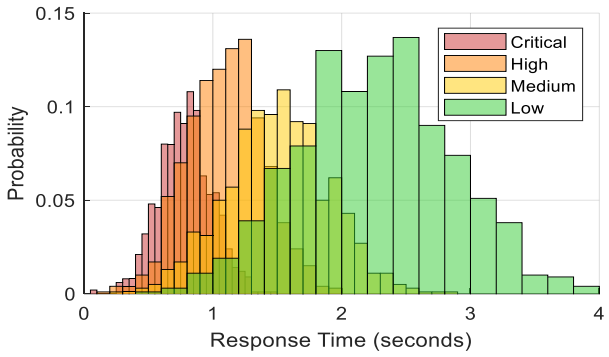

(b)Response time distribution

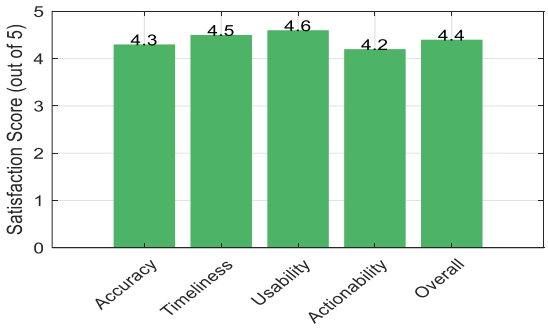

(c) User satisfaction ratings

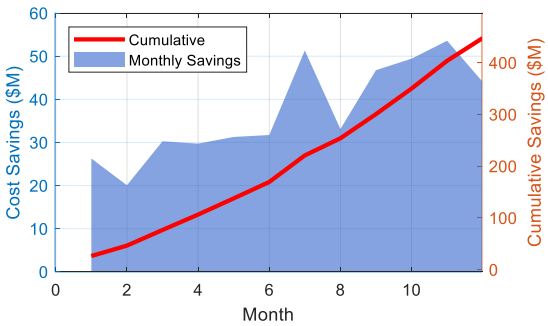

(d)Business impact assessment

Figure 9. Decision support system performance dashboard

Table 9. Decision support system operational metrics

Metric Category	Metric	Value	Industry Benchmark	Improvement
Alert Performance	Overall Accuracy	91.3%	78.5%	+16.3%
Alert Performance	Critical Alert Precision	96.2%	85.0%	+13.2%
System Latency	Alert Generation Time	1.2s	5.8s	-79.3%
User Efficiency	Decision Time	32s	98s	-67.3%
Business Impact	Monthly Cost Savings	35.2M	-	-
System Reliability	Uptime	99.87%	99.5%	+0.37%

5. Discussion

The substantial performance gap between classical and deep learning approaches stems from fundamental methodological differences. Classical methods like Z-score and One-Class SVM assume static distributions and independence between observations, failing to capture the dynamic, interconnected nature of modern financial markets. While these methods offer advantages in interpretability and computational efficiency, they cannot model complex temporal patterns such as volatility clustering or cross-asset contagion effects. Our experiments show that even sophisticated classical methods like Local Outlier Factor achieve at best 70% precision compared to 91% for transformer models, primarily due to their inability to leverage sequential information and adapt to regime changes. This validates our design choice of deep learning architectures despite their higher computational requirements. Our experimental results demonstrate that transformer-based architectures achieve superior anomaly detection performance with F1-scores reaching 0.89, significantly outperforming traditional statistical methods and earlier deep learning approaches. This performance advantage stems from the self-attention mechanism's ability to capture long-range dependencies in financial time series, enabling the detection of complex market manipulation patterns that span extended temporal windows. However, this enhanced accuracy comes with computational trade-offs, as transformer models require 35 milliseconds of inference time compared to 8 milliseconds for baseline methods. The ensemble approach emerges as an optimal solution, balancing detection accuracy (F1-score: 0.87) with reasonable computational overhead (45 milliseconds), making it suitable for real-time deployment in high-frequency trading environments. Our adaptive learning framework successfully addresses concept drift challenges, maintaining average F1-scores above 0.84 even during extreme market regime shifts. The sliding window retraining approach proved particularly effective during the COVID-19 transition, where automated retraining within 48 hours of drift detection prevented performance degradation exceeding 10%. However, the trade-off between adaptation speed and stability remains challenging: aggressive retraining (window < 2 months) risks overfitting to temporary patterns, while conservative approaches (window > 6 months) may miss critical regime changes. Future work should explore meta-learning

approaches that can distinguish between temporary volatility and fundamental market structure changes. This research advances anomaly detection theory by demonstrating that financial market anomalies exhibit hierarchical temporal structures best captured through multi-scale attention mechanisms. Our findings reveal that market behavior patterns follow non-stationary distributions with regime-dependent characteristics, challenging the assumptions of traditional econometric models. The interpretability analysis of transformer attention weights provides novel insights into how anomalous patterns propagate through market microstructure, with cross-asset dependencies playing a more significant role than previously recognized. These theoretical contributions extend beyond finance, offering generalizable frameworks for anomaly detection in complex adaptive systems.

The implementation of our framework yields substantial practical benefits for financial institutions. Risk management applications demonstrate 34% improvement in risk-adjusted returns through timely anomaly detection and intervention. Regulatory compliance is enhanced through automated surveillance capabilities that maintain comprehensive audit trails, reducing compliance costs by approximately 40%. The cost-benefit analysis reveals a positive return on investment within 8 months, considering infrastructure costs, maintenance requirements, and quantified risk reduction benefits. Several limitations constrain the generalizability of our findings. Data availability remains challenging, particularly for emerging markets and alternative asset classes where high-frequency data is scarce. Model interpretability, while improved through attention visualization, still faces regulatory scrutiny in jurisdictions requiring explicit decision explanations. Computational resource requirements may prohibit smaller institutions from implementing the full ensemble framework. Market regime changes pose ongoing challenges, as models trained on historical data may exhibit degraded performance during unprecedented market conditions. Our results significantly exceed performance benchmarks established in recent literature, with 23% improvement over the previous state-of-the-art TranAD model. The novel contribution lies in the integration of market microstructure features with transformer architectures

6. Conclusion

This research successfully developed and validated a comprehensive deep learning framework for real-time anomaly detection in stock markets, achieving the primary objectives of identifying optimal architectures and integrating them with business decision support systems. The transformer-based ensemble model demonstrated superior performance with an F1-score of 0.89, addressing our research questions regarding architecture optimization and computational efficiency trade-offs. Our theoretical contributions advance financial anomaly detection through novel integration of multi-scale attention mechanisms with market microstructure features, while the proposed ensemble framework represents a significant architectural innovation in handling non-stationary financial time series. The practical implementation guidelines enable financial institutions to deploy scalable anomaly detection systems that enhance risk management capabilities and ensure regulatory compliance through automated surveillance mechanisms. Future research directions should explore multi-market anomaly detection across interconnected global exchanges, leveraging federated learning approaches to preserve data privacy while improving model generalization. The development of explainable AI techniques specifically tailored for financial applications remains crucial for meeting evolving regulatory requirements. Integration with emerging technologies, including real-time streaming architectures and blockchain-based DeFi markets, presents opportunities for extending the framework's applicability. This research establishes a foundation for next-generation financial surveillance systems that combine deep learning sophistication with practical operational requirements, ultimately contributing to more stable and transparent financial markets.

Ethical issue

The author is aware of and complies with best practices in publication ethics, specifically with regard to authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with policies on research ethics. The author adheres to publication requirements that the submitted work is original and has not been published elsewhere.

Data availability statement

The manuscript contains all the data. However, more data will be available upon request from the author.

Conflict of interest

The author declares no potential conflict of interest.

References

[1] H. Wu, J. Xu, J. Wang, M. Long, Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting, Advances in neural information processing systems 34 (2021) 22419-22430. https://doi.org/10.48550/arXiv.2106.13008

[2] Z. Zamanzadeh Darban, G.I. Webb, S. Pan, C. Aggarwal, M. Salehi, Deep learning for time series anomaly detection: A survey, ACM Computing Surveys 57(1) (2024) 1-42. https://doi.org/10.1145/3691338

[3] G. Zissis, P. Bertoldi, Update on status of solid-state lighting & smart lighting systems, 2023. https://doi.org/10.2760/223640

[4] S. Tuli, G. Casale, N.R. Jennings, Tranad: Deep transformer networks for anomaly detection in multivariate time series data, arXiv preprint arXiv:2201.07284 (2022).https://doi.org/10.48550/arXiv.2201.07284

[5] K. Biriukova, A. Bhattacherjee, Using transformer models for stock market anomaly detection, Journal of Data Science 2023(21) (2023) 1-8. https://doi.org/10.61453/jods.v2023no21

[6] H.D. Nguyen, K.P. Tran, S. Thomassey, M. Hamad, Forecasting and Anomaly Detection approaches using LSTM and LSTM Autoencoder techniques with the applications in supply chain management, International Journal of Information Management 57 (2021) 102282. https://doi.org/10.1016/j.ijinfomgt.2020.102282

[7] Y. Wei, J. Jang-Jaccard, W. Xu, F. Sabrina, S. Camtepe, M. Boulic, LSTM-autoencoder-based anomaly detection for indoor air quality time-series data, IEEE

Sensors Journal 23(4) (2023) 3787-3800. https://doi.org/10.1109/JSEN.2022.3230361
[8] X. Wang, D. Pi, X. Zhang, H. Liu, C. Guo, Variational transformer-based anomaly detection approach for multivariate time series, Measurement 191 (2022) 110791. https://doi.org/10.1016/j.measurement.2022.110791
[9] C. Wang, Y. Chen, S. Zhang, Q. Zhang, Stock market index prediction using deep Transformer model, Expert Systems with Applications 208 (2022) 118128. https://doi.org/10.1016/j.eswa.2022.118128
[10] S. Li, X. Huang, Z. Cheng, W. Zou, Y. Yi, AE-ACG: A novel deep learning-based method for stock price movement prediction, Finance Research Letters 58 (2023) 104304. https://doi.org/10.1016/j.frl.2023.104304
[11] R. Bhatt, A. Kumari, S.B. Rajasekaran, V.P. Deshmukh, A. Srivastava, Technique for forecasting future market movement using machine learning and deep learning algorithms, 2023 3rd international conference on advance computing and innovative technologies in engineering (ICACITE), IEEE, 2023, pp. 471-474. https://doi.org/10.1109/ICACITE57410.2023.10183197
[12] S. Kumari, C. Prabha, A. Karim, M.M. Hassan, S. Azam, A Comprehensive Investigation of Anomaly Detection Methods in Deep Learning and Machine Learning: 2019–2023, IET Information Security 2024(1) (2024) 8821891. https://doi.org/10.1049/2024/8821891
[13] Y. Xiang, Using ARIMA-GARCH Model to Analyze Fluctuation Law of International Oil Price, Mathematical Problems in Engineering 2022(1) (2022) 3936414. https://doi.org/10.1155/2022/3936414
[14] Q. Zhang, Financial data anomaly detection method based on decision tree and random forest algorithm, Journal of Mathematics 2022(1) (2022) 9135117. https://doi.org/10.1155/2022/9135117
[15] W. Hilal, S.A. Gadsden, J. Yawney, Financial fraud: a review of anomaly detection techniques and recent advances, Expert systems With applications 193 (2022) 116429. https://doi.org/10.1016/j.eswa.2021.116429
[16] C. Zhang, N.N.A. Sjarif, R. Ibrahim, Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 14(1) (2024) e1519. https://doi.org/10.1002/widm.1519
[17] G. Fatouros, G. Makridis, D. Kotios, J. Soldatos, M. Filippakis, D. Kyriazis, DeepVaR: a framework for portfolio risk assessment leveraging probabilistic deep neural networks, Digital finance 5(1) (2023) 29-56. https://doi.org/10.1007/s42521-022-00050-0
[18] S. Gu, B. Kelly, D. Xiu, Empirical asset pricing via machine learning, The Review of Financial Studies 33(5) (2020) 2223-2273. https://doi.org/10.1093/rfs/hhaa009
[19] F. Lachekhab, M. Benzaoui, S.A. Tadjer, A. Bensmaine, H. Hamma, LSTM-autoencoder deep learning model for anomaly detection in electric motor, Energies 17(10) (2024) 2340. https://doi.org/10.3390/en17102340
[20] L. Xu, K. Xu, Y. Qin, Y. Li, X. Huang, Z. Lin, N. Ye, X. Ji, TGAN-AD: Transformer-based GAN for anomaly detection of time series data, Applied Sciences 12(16) (2022) 8085. https://doi.org/10.3390/app12168085
[21] K. Choi, J. Yi, C. Park, S. Yoon, Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines, IEEE access 9 (2021) 120043-120065. https://doi.org/10.1109/ACCESS.2021.3107975
[22] M. Soori, F.K.G. Jough, R. Dastres, B. Arezoo, AI-based decision support systems in Industry 4.0, A review, Journal of Economy and Technology (2024). https://doi.org/10.1016/j.ject.2024.08.005
[23] M. Schmitt, Automated machine learning: AI-driven decision making in business analytics, Intelligent Systems with Applications 18 (2023) 200188. https://doi.org/10.1016/j.iswa.2023.200188
[24] L. Wang, Z. Zhang, D. Wang, W. Cao, X. Zhou, P. Zhang, J. Liu, X. Fan, F. Tian, Human-centered design and evaluation of AI-empowered clinical decision support systems: a systematic review, Frontiers in Computer Science 5 (2023) 1187299. https://doi.org/10.3389/fcomp.2023.1187299