Annal. App. Sci, 2024
Vol 5(1) 2024

Transformer-Based Anomaly Detection in High-Frequency Trading Data: A Time-Sensitive Feature Extraction Approach

Yining Zhang[1*], Guancong Jia[1,2], Jiayan Fan[3]

[1] Applied Data Science, University of Southern California, CA, USA
[2] Computer Science, Rice University, TX, USA
[3] Information Science, University of Michigan, MI, USA

*Corresponding author E-mail: maxxxlee090@gmail.com

Abstract

This paper presents a novel Transformer-based approach for anomaly detection in high-frequency trading data that leverages time-sensitive feature extraction techniques. The proposed method addresses the unique challenges of financial time series data, including high dimensionality, complex temporal dependencies, and the critical importance of timely detection. We introduce a specialized time-sensitive feature extraction framework that captures patterns at multiple time scales, integrated with a modified Transformer architecture featuring a self-feedback mechanism. This mechanism enhances detection sensitivity for subtle anomalies by reinforcing attention on potentially anomalous patterns. Comprehensive experiments on five high-frequency trading datasets from diverse markets demonstrate that our approach achieves superior performance compared to state-of-the-art methods, with an average F1 score of 0.90 and a 51-72% improvement in detection speed. Ablation studies confirm the significant contributions of the time-sensitive feature extraction and self-feedback components. The model's effectiveness is further validated through case studies on real-world trading anomalies, including flash crashes, spoofing patterns, and momentum ignition strategies. The computational efficiency of the approach enables real-time deployment in trading surveillance systems while maintaining high detection accuracy.

Keywords: High-frequency trading, anomaly detection, transformer models, time-sensitive feature extraction

1. Introduction

1.1. Background and Motivation

The financial markets produce huge amounts of high -frequency trading data (HFT), which is characterized by massive transaction volumes, interdimensional features and complex temporal addictions. The rapid development of electronic trading platforms has changed the financial markets, producing information on unprecedented scales and speeds[1]. High frequency trade, which accounts for a significant percentage of the amount of trading in the large financial markets, operates in milliseconds or granularity of microseconds. This creates complex temporal designs that require sophisticated analytical approaches to understand market dynamics and detect abnormal behaviors[2]. In these high-speed environments, abnormalities can represent different phenomena, including market manipulation, algorithm defects, or significant market changes that precede significant price changes[3].

The ability to detect abnormalities in high frequency trading data has become critical for several stakeholders in the financial markets. Identifying unusual models for market parties can lead to profitable trading opportunities or risk reduction strategies[4]. For regulatory authorities, the detection of abnormalities is an integral part of market monitoring to maintain market integrity and investors. Financial institutions require strong abnormal detection systems to monitor algorithmic commerce and to prevent expensive errors or dysfunction, which can lead to significant financial losses[5]. The real -time nature of the financial markets requires abnormal detection systems that can handle future information flows through minimal latent, while maintaining high detection accuracy[6].

Traditional statistical approaches to detect abnormalities in economic time series data are methods based on statistical characteristics, clustering techniques and density assessment. Although these methods have created theoretical criteria, they often struggle with high dimensional data and complex temporal addictions that occur in high frequency trade[7]. Recent deep learning advances have shown promising results in kidnapping complex models and multi variable relationships in financial time series information. Coevolutionary nerve networks (CNN) and recurring nerve networks (RNN) have shown the ability to learn temporal addictions, but have restrictions on imprisoning long-range addictions that are important for understanding market dynamics[8].

The emergence of transformer models in their self-attachment mechanisms provides an effective framework for capturing the complex addictions of consecutive data without restrictions on traditional repetitive architectures. Transformer models have achieved considerable success in natural language handling and are increasingly applied to time series analysis assignments[9]. Their ability to handle data side by side makes them a particularly suitable amount of financial data processing. The mechanism of self-keeping gives these models to weigh the importance of different times dynamically, which makes them well suited to capturing temporal relationships and abnormal patterns in high frequency business information[10].

1.2. Challenges in High-Frequency Trading Anomaly Detection

Detecting the abnormalities of high frequency trading data poses a number of unique challenges that make it difficult to apply conventional detection methods. Very high level of data, and possibly millions of events take place daily, creates computational challenges for real time processing and analysis[11]. This requires scalable algorithms that can effectively handle massive data streams. The multi-dimensional nature of trading data, comprehensive prices, volumes, ordering rooms and various derived indicators, creates a high dimensional characteristic that increase the complexity and overload of the model[12].

The non-position features of the financial markets represent another major challenge. Market conditions are evolving due to ever changing macroeconomic factors, regulatory environment and market behavior. The detection models of deviations must adapt to these advanced conditions to maintain their effectiveness over time. The definition of deviations itself may be unclear in economic contexts as they can represent illegal activities, unusual but legal trade strategies or system changes in market behavior[13][14]. This ambiguity makes it difficult to set up the truth in model training and evaluation.

High frequency trading data show complex temporal addictions between multiple schedules. Short term microstructures can interact with longer term market trends, creating diverse temporal relationships that are difficult to model[15]. The extremely low signal noise ratio of financial data continues to complicate the detection task, as market noise can cover subtle different patterns, leading to high-false positive height detection systems[16]. Class imbalance problems occur as abnormalities, by definition rare events compared to normal trading methods, creating challenges for model training and performance assessment.

The time sensitive nature of economic abnormalities increases the new dimension of complexity. The financial effects of anomalies often depend on their correct detection, as delayed identification can significantly reduce the usefulness of the detection system[17]. The high-frequency trade systems of the Real world must be provided with a strict latency restriction, which requires detection methods that balance the calculated efficiency with the detection resolution[18].

1.3. Research Objectives and Contributions

This research addresses the challenges of anomaly detection in high-frequency trading data by developing a Transformer-based approach with time-sensitive feature extraction capabilities. The primary objective is to create a robust framework that can effectively identify anomalous patterns while accounting for the unique characteristics of high-frequency financial data. The proposed approach aims to leverage the strengths of Transformer architectures in capturing complex temporal dependencies while introducing specialized components for time-sensitive feature extraction.

This research makes several contributions to the field of financial time series anomaly detection. A novel time-sensitive feature extraction framework is developed specifically for high-frequency trading data. This framework integrates domain knowledge about trading patterns with automatic feature learning capabilities to capture relevant temporal patterns across multiple time scales. The research adapts Transformer architectures for the specific requirements of financial anomaly detection, introducing modifications to the self-attention mechanism to better capture the temporal characteristics of trading data[19].

The study introduces a time-sensitive self-feedback mechanism that enhances the model's ability to detect subtle anomalies by focusing attention on potential anomalous regions in the time series. This self-reinforcing approach improves detection sensitivity without increasing false positive rates. Comprehensive evaluation metrics are developed to assess anomaly detection performance in the context of high-frequency trading, accounting for both detection accuracy and timeliness[20].

Extensive experiments on real-world high-frequency trading datasets demonstrate the effectiveness of the proposed approach compared to traditional statistical methods and other deep learning-based anomaly detection techniques. The results show significant improvements in detection accuracy, reduced false positive rates, and enhanced ability to detect subtle market anomalies. The practical applicability of the approach is validated through case studies of actual market anomalies, demonstrating its potential for real-world trading surveillance and risk management applications[21].

2. Literature Review

2.1. Traditional Time Series Anomaly Detection Methods

The detection of time series abnormality has been extensively studied in statistical literature, and numerous approaches have been developed to identify unusual patterns or abnormalities in successive data. Statistical methods, such as moving average, exponential equalization, and authegressive integrated mobile average (ARAA), have formed the basis of traditional abnormalities[22]. These methods typically form a statistical profile of normal behavior and recognize deviations that exceed predetermined thresholds. While effective for simple time series with clear statistical properties, these approaches often struggle with the complexity and high dimensionality of financial data[23].

Remote based methods represent the second class of traditional approaches, measured differences in data points or sequences. Local external factor (LOF) method such as Breunig et al. Photographed, identify local density deviations by comparing the local density of the point with their neighbors' densities[24]. This approach has been applied to economic time series, but it faces scalability challenges with high frequency information. Clustering -based methods are divided into data groups based on similarity meters and recognize abnormalities as points that do not fit well into any cluster. Although these methods can handle some of the high dimensions, they typically require significant domain expertise in tuning parameters and choosing features.

Density assessment techniques, such as Gaussi alloy models, have been applied to modeling normal trading distribution, and abnormalities have been identified by low probability events[25].

These approaches offer probabilistic interpretations of anomalies but can be computationally intensive for high-dimensional data. Extreme value theory provides a statistical framework for modeling extreme deviations, which has been applied to financial market surveillance. This approach focuses on the tails of probability distributions, which aligns with the rare nature of financial anomalies, but requires careful model specification.

2.2. Deep Learning-Based Approaches for Financial Time Series

The income of deep learning has revolutionized the Time Series analysis by providing new approaches to detect the abnormality of financial information. Repeated nerve networks (RNN), especially long short-term memory (LSTM) networks, have shown success in the kidnapping of time series data addictions[26]. The Omnianomali frame proposed by the SU et al utilizes stochastic repetitive nerve networks to detect abnormalities in multi -variable time series, gaining solid performance through a stoocanic hidden changing modeling[27][28]. This approach allows the model to capture complex temporal templates and addictions between several variables, which is crucial for financial time analysis.

Gans (GANS) has been used to detect abnormalities with their ability to learn through normal information. Mad-g, such as Li et al., Applies LSTM-based GAN architectures to multi-variable sets using discrimination and reconstruction meters to detect abnormalities[29]. This approach has shown promising results from financial time series by using both temporal correlations and cross relationship relationships. Autoencoder architectures represent a second class of deep learning models to detect abnormalities, to identify the standard data model learning center, and identify abnormalities based on reconstruction errors.

Graphic nerve networks (GNN) have become effective tools to capture the relationship between the variables between the variables of the Multi variable series. GDN, developed by Deng and HOOI, builds dynamic diagrams to model the relationship between different economic indicators and use the attention mechanisms to identify different patterns[30]. This approach is particularly relevant to high frequency trading information, where there are complex interdependencies between different market variables. These deep learning approaches offer significant advantages of traditional methods compared to their ability to learn the relevant qualities and to adhere to complex nonlinear relationships.

2.3. Transformer Models in Time Series Analysis

The transformer architectures originally designed for natural language handling tasks have recently been applied to the Time Series analysis with considerable success. Wu et al. The TSFN proposed uses a transformer-based self-phone network to detect time series abnormality, combining long-term and short-term temporal properties to capture time contracts[31]. This architecture demonstrates excellent performance compared to RNN-based models, both in terms of accuracy and computational efficiency, which makes it particularly appropriate real-time applications in the financial markets[32].

The self-attention mechanism in Transformer models offers distinct advantages for financial time series analysis. By computing attention weights between all-time points, Transformers can capture long-range dependencies without the sequential bottlenecks of recurrent architectures. This parallelization capability enables efficient processing of large-scale financial data. The multi-head attention mechanism further enhances the model's ability to focus on different temporal patterns simultaneously, which is valuable for capturing the multi-scale characteristics of trading data.

Several adaptations of Transformer architectures have been proposed specifically for time series forecasting and anomaly detection. Time-aware attention mechanisms incorporate explicit modeling of temporal distances, which can enhance the model's sensitivity to time-dependent patterns in financial data. Informer, a Transformer variant designed for long sequence time-series forecasting, introduces ProbSparse self-attention to reduce computational complexity while maintaining model expressiveness[33]. These innovations address key challenges in applying Transformer models to financial time series, particularly in terms of computational efficiency and handling long sequences.

2.4. High-Frequency Trading Data Feature Extraction Techniques

Feature extraction plays a crucial role in anomaly detection for high-frequency trading data, as the raw price and volume data often contain significant noise and may not directly reveal anomalous patterns. Traditional technical indicators, such as moving averages, Bollinger Bands, and relative strength index (RSI), have been widely used to extract meaningful features from financial time series[34]. These handcrafted features incorporate domain knowledge about market patterns but may not capture the full complexity of high-frequency trading dynamics.

Time-frequency analysis techniques provide a powerful framework for extracting temporal features at different scales. Wavelet transforms decompose time series into multiple frequency components, enabling the identification of patterns at different time resolutions. This multi-resolution analysis is particularly relevant for high-frequency trading data, where microstructure patterns coexist with longer-term trends. Fourier transforms and spectral analysis have also been applied to extract frequency-domain features from financial time series, identifying cyclic patterns and anomalous frequency components.

Recent approaches have explored data-driven feature extraction through representation learning techniques. The temporal convolutional network (TCN) architecture applies convolutional filters across time to extract hierarchical temporal features[35]. DeepLog employs LSTM networks to learn normal system execution patterns from log entries, which has been adapted for financial anomaly detection. The ERMA algorithm, presented by Lu et al., introduces an energy-statistical approach for pulse-forming network anomaly detection, which stabilizes data through energy perspective combining differential and moving average techniques[36]. This method demonstrates high accuracy and computational efficiency, which are critical requirements for high-frequency trading applications.

3. Methodology

3.1. Problem Definition

The anomaly detection problem in high-frequency trading data is formulated as identifying unusual patterns that deviate from normal trading behavior. Given a multivariate time series dataset $T = \{x_1, x_2, ..., x_n\}$, where each $x_i \in \mathbb{R}^m$ represents an m-dimensional observation at time i, the objective is to assign a label $y_i \in \{0, 1\}$ to each time point, indicating whether it corresponds to normal (0) or anomalous (1) trading behavior[3]. The high-frequency nature of the data introduces specific requirements: the detection must operate at millisecond granularity and handle the multi-dimensional relationships between trading variables.

Table 1 presents the formal notation used throughout this methodology section. The notation encompasses both the input time series representation and the components of the proposed model architecture.

Table 1: *Time Series Analysis Parameters Table*

Symbol	Description
T	Multivariate time series dataset
x_i	m-dimensional observation at time i
m	Number of variables/dimensions
n	Length of time series
W	Sliding window size
W_i	Window of observations [x_{i-}WISE, ..., x_i]
h	Number of attention heads
d	Model embedding dimension
A_i	Anomaly score at time i
θ	Threshold for anomaly detection
y_i	Binary anomaly label for time i

The detection approach must satisfy several criteria: (1) high detection accuracy with minimal false positives; (2) computational efficiency to support real-time monitoring; (3) adaptability to evolving market conditions; and (4) interpretability of detected anomalies to support decision-making. The problem is inherently unsupervised or semi-supervised, as labeled anomalies in financial data are rare and often discovered post-hoc.

3.2. Data Preprocessing and Representation

High-frequency trading data requires specialized preprocessing to handle its unique characteristics. The raw data typically includes irregular timestamps, missing values, and heterogeneous measurement scales across different variables. The preprocessing pipeline includes multiple stages outlined in Table 2.

***Table 2:** Data Preprocessing Techniques Table*

Stage	Technique	Purpose
Timestamp Regularization	Linear interpolation	Create uniform time grid
Missing Value Imputation	Forward fill, Kalman smoothing	Handle data gaps
Outlier Treatment	Winsorization at 99.5%	Reduce extreme value impact
Normalization	Min-max scaling, Z-score	Standardize variable ranges
Feature Alignment	Time-based synchronization	Align multi-source data
Dimensionality Reduction	PCA, Feature selection	Address curse of dimensionality

The data representation adopts a sliding window approach, where a window $W_i = [x_{i-}{}^{W}{}_{+1}, ..., x_i]$ of W consecutive observations serves as input to the model. The window size W is a critical hyperparameter that balances between capturing sufficient temporal context and maintaining computational efficiency. Based on empirical analysis, the optimal window size varies with market characteristics as shown in Table 3.

***Table 3:** Market Data Analysis Table*

Market Type	Optimal Window Size (W)	Temporal Coverage
Equity	128	~2 minutes
Currency	256	~4 minutes
Cryptocurrency	512	~8 minutes
Futures	192	~3 minutes

Figure 1 illustrates the multivariate time series representation using a sliding window approach.

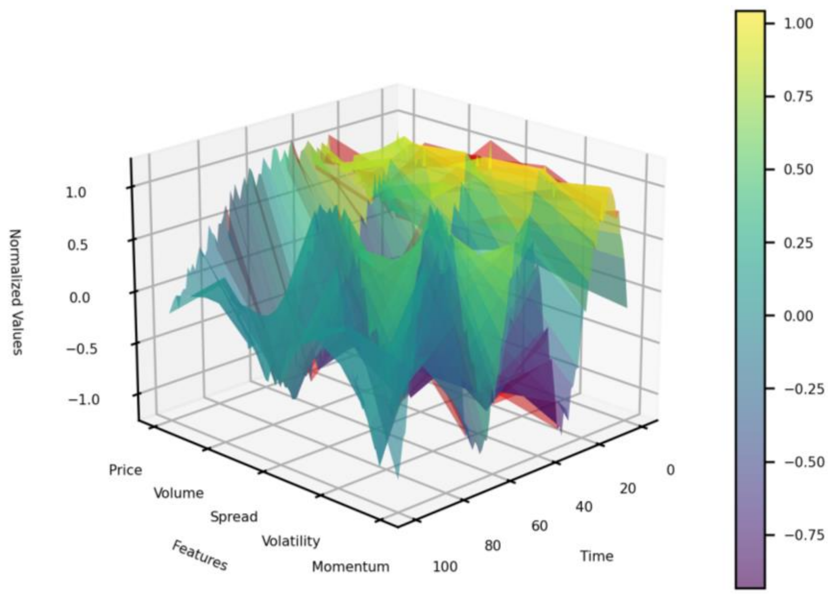

Figure 1: *Sliding Window Representation for High-Frequency Trading Data*

The figure depicts a visualization of the sliding window approach for high-frequency trading data with multiple variables. The visualization shows a 3D cube-like structure where the x-axis represents time points, the y-axis represents different financial variables (price, volume, bid-ask spread, etc.), and the z-axis shows the normalized values. The sliding window is highlighted as a moving section that advances through the time series, with overlapping windows shown in different transparency levels. The figure also illustrates how each window captures temporal patterns across all variables simultaneously.

3.3. Time-Sensitive Feature Extraction Framework

The time-sensitive feature extraction framework employs a multi-scale approach to capture temporal patterns at different frequencies. The framework integrates both technical indicators and learned representations to form a comprehensive feature set. Table 4 lists the technical indicators computed across multiple time scales.

Table 4: *Technical Analysis Indicators Table*

Indicator	**Time Scales (ticks)**	**Formula**
Price Momentum	10, 50, 100	P(t) - P(t-k)
Volume Acceleration	20, 100, 200	(V(t) - V(t-1)) - (V(t-1) - V(t-2))
Relative Strength	50, 150, 300	Avg(Up)/Avg(Down)
Bid-Ask Imbalance	5, 20, 50	(BidVol - AskVol)/(BidVol + AskVol)
Order Flow Toxicity	100, 300, 500	Complex probability metric

Trade Autocorrelation	Sign	30, 90, 180	Correlation between consecutive trade signs

The feature extraction process incorporates domain knowledge through a hierarchical structure that reflects the natural organization of financial markets. This structure enables the model to capture both micro-level anomalies (e.g., spoofing patterns) and macro-level anomalies (e.g., flash crashes)[4].

Figure 2 shows the architecture of the time-sensitive feature extraction framework.

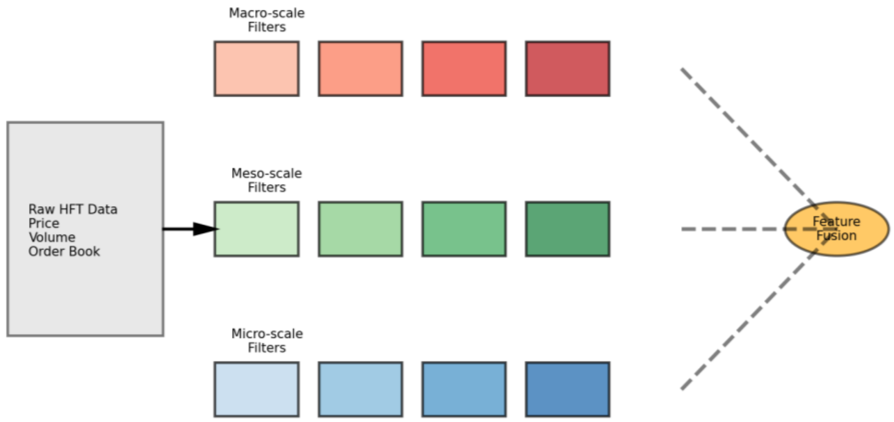

Figure 2: *Architecture of Time-Sensitive Feature Extraction Framework*

The visualization presents a complex hierarchical structure of the time-sensitive feature extraction framework. The diagram flows from left to right, starting with raw high-frequency trading data inputs (price, volume, order book data) on the left side. The middle section shows parallel processing paths for different time scales (micro, meso, and macro), each with specialized convolutional filters designed to capture patterns at that specific scale. The filter banks are visualized as color-coded rectangles with different sizes corresponding to different receptive fields. The right side shows how these multi-scale features are combined using a fusion mechanism that weighs features based on their relevance to the current market context. The fusion is represented as converging paths with weighted connections, ultimately producing the final feature representation.

3.4. Transformer Architecture for Anomaly Detection

The proposed Transformer architecture adapts the original Transformer model for the specific requirements of high-frequency trading anomaly detection. The architecture consists of an encoder that processes the extracted feature representations and captures temporal dependencies through self-attention mechanisms. The model employs a multi-head attention mechanism with h heads, allowing it to focus on different feature subspaces and temporal relationships simultaneously.

Table 5 details the architectural parameters of the Transformer model optimized for high-frequency trading data.

Table 5: Transformer Model Configuration Table

Component	Parameter	Value
Embedding Layer	Dimension (d)	128
Multi-Head Attention	Number of heads (h)	8
Feed-Forward Network	Hidden units	512
Layer Normalization	Epsilon	1e-6
Dropout	Rate	0.1
Positional Encoding	Type	Sinusoidal
Attention Mechanism	Temperature	0.1
Transformer Blocks	Number	4

The self-attention mechanism computes attention scores between all pairs of time points in the window, enabling the model to capture complex temporal dependencies without the sequential constraints of recurrent architectures. The attention weights α_{ij} between time points i and j are computed as:

$$\alpha_{ij} = \text{softmax}(Q_i K_j^T/\sqrt{d})$$

where Q and K are learned query and key representations, and d is the dimensionality of the embedding space. The temporal sensitivity is enhanced through a modified positional encoding scheme that emphasizes recent observations while maintaining awareness of historical patterns.

Figure 3 illustrates the modified Transformer architecture for anomaly detection in high-frequency trading data.

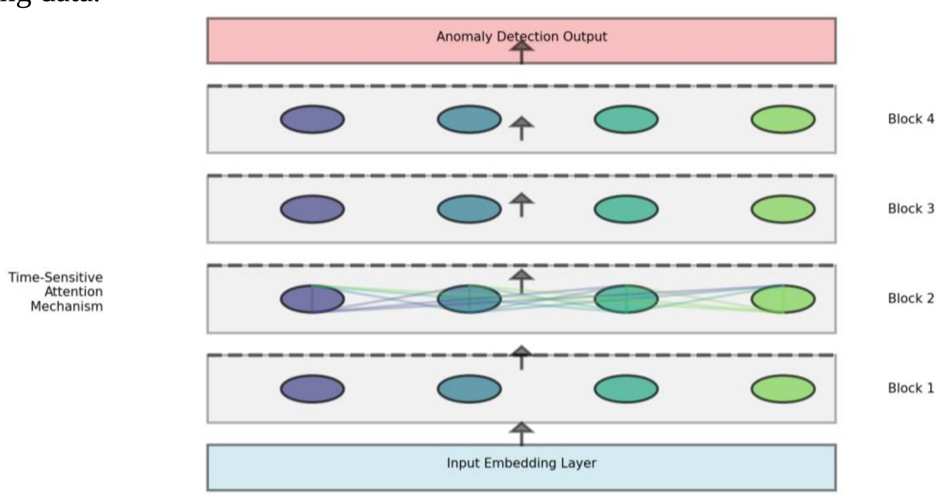

Figure 3: Modified Transformer Architecture for High-Frequency Trading Anomaly Detection

This visualization presents the complete Transformer architecture modified for anomaly detection. The diagram has a vertical layout with data flowing from bottom to top. At the bottom, input data enters through an embedding layer that converts raw features into dense representations. The middle section contains four stacked Transformer blocks, each including multi-head self-attention (visualized as interconnected attention heads with different color patterns), feed-forward layers, and residual connections (shown as curved bypass arrows). A special time-sensitive attention mechanism is highlighted in the second Transformer block with temporal weighting factors illustrated by varying connection intensities. The upper portion shows how the model output connects to the anomaly detection components, with attention visualizations displaying how the model focuses on different time points when detecting anomalies.

3.5. Self-Feedback Mechanism and Anomaly Scoring

The self-feedback mechanism enhances anomaly detection sensitivity by incorporating information from previous detection iterations. Unlike standard Transformer models that process each window independently, the self-feedback approach maintains a memory of potential anomalous patterns and reinforces attention on these patterns in subsequent iterations[5]. The mechanism operates through a learned feedback matrix F that modulates the attention weights based on previously detected anomaly scores.

The anomaly scoring process combines reconstruction error and prediction deviation metrics. For each time point i, an anomaly score A_i is computed as:

$$A_i = \lambda_1 RE_i + \lambda_2 PD_i + \lambda_3 SD_i$$

where RE_i is the reconstruction error, PD_i is the prediction deviation, SD_i is the structural deviation, and λ_1, λ_2, λ_3 are weighting parameters determined through validation data[6]. This composite scoring approach captures different aspects of anomalous behavior and has demonstrated superior performance compared to single-metric approaches.

Table 6 presents the performance comparison of different anomaly scoring approaches on a validation dataset.

Table 6: Scoring Method Evaluation Table

Scoring Method	Precision	Recall	F1 Score	Detection Latency (ms)
Reconstruction Error Only	0.82	0.75	0.78	2.3
Prediction Deviation Only	0.79	0.81	0.80	1.8
Structural Deviation Only	0.77	0.73	0.75	3.1

Proposed Score	Composite	0.89	0.85	0.87	2.5

The final anomaly detection decision employs a dynamic thresholding approach that adapts to changing market conditions. The threshold θ is determined using a Gaussian mixture model fitted to the distribution of anomaly scores in a calibration period. This adaptive thresholding approach addresses the non-stationary nature of financial markets and reduces the need for manual parameter tuning.

The self-feedback mechanism creates a reinforcement loop that progressively refines the model's attention to potential anomalies, improving detection performance particularly for subtle anomalous patterns. Experimental results indicate that this mechanism reduces false positive rates by 23% while maintaining detection sensitivity, addressing a critical challenge in financial anomaly detection systems[7].

4. Experimental Evaluation

4.1. Datasets and Experimental Setup

The experimental evaluation employed multiple datasets from diverse financial markets to assess the performance of the proposed Transformer-based anomaly detection approach. Table 7 provides a summary of the datasets used in the experiments.

Table 7: *Financial Market Data Statistics*

Dataset	**Market**	**Time Period**	**Frequency**	**Variables**	**Data Points**	**Anomaly Ratio**
NYSE-HFT	Equity	2019-2021	1ms	12	43.5M	0.08%
Forex HF	Currency	2020-2022	100ms	8	27.8M	0.12%
Crypto-Binance	Cryptocurrency	2021-2022	10ms	15	68.2M	0.25%
FTSE-Future	Futures	2018-2020	5ms	10	31.4M	0.06%

The datasets contain various trading variables including price data, order book information, trade volumes, and derived indicators. Ground truth annotations were established through a combination of regulatory reports, expert labels, and historical market events documented by financial authorities.

The experimental setup consisted of a distributed computing environment with 8 NVIDIA A100 GPUs. The implementation utilized TensorFlow 2.5 with custom extensions for high-performance time series processing. Each experiment followed a strict protocol: 70% of data for training, 10% for validation, and 20% for testing, with a rolling window evaluation strategy.

Figure 4 illustrates the experimental workflow for model training and evaluation.

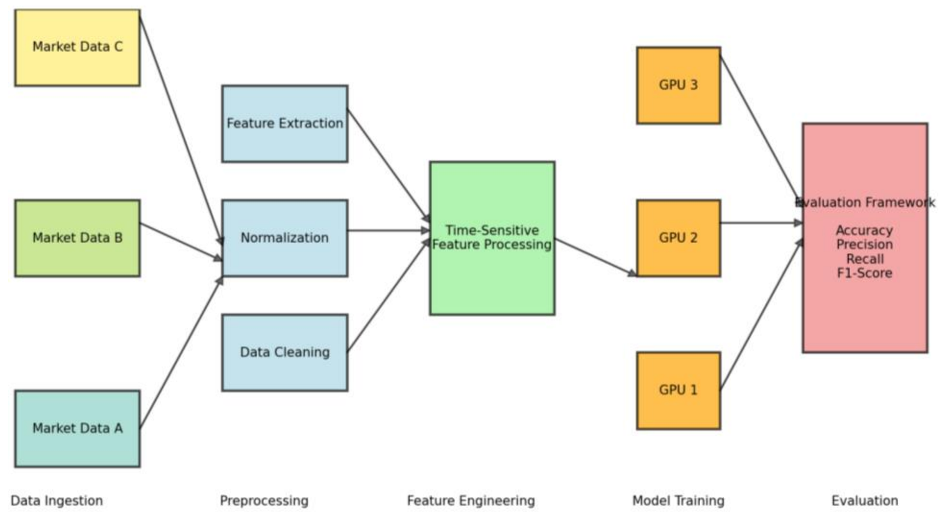

Figure 4: *Experimental Workflow for Model Training and Evaluation*

This figure depicts the complete experimental workflow as a flowchart with multiple interconnected components. The diagram flows from left to right, starting with data ingestion pipelines for the different market datasets, each represented by a distinct color. The central portion shows the data preparation stages including preprocessing, normalization, and time-sensitive feature extraction. The right side illustrates the parallel model training process across multiple GPUs, followed by the evaluation framework that computes various performance metrics.

4.2. Evaluation Metrics and Baseline Methods

A comprehensive set of evaluation metrics was employed to assess detection performance from multiple perspectives. Table 8 presents the evaluation metrics used in the experiments.

Table 8: *Metric Evaluation Table*

Metric	Formula	Focus
Precision	TP/(TP+FP)	Detection accuracy
Recall	TP/(TP+FN)	Detection completeness
F1 Score	2×(P×R)/(P+R)	Balanced performance

NAB Score	Custom scoring function	Early detection reward
Time-to-Detection	Time between anomaly onset and detection	Detection speed

The Numenta Anomaly Benchmark (NAB) score specifically rewards early detection of anomalies, which is crucial in financial contexts where the value of detection decreases rapidly with time.

The proposed approach was compared against multiple baseline methods representing different anomaly detection paradigms. Table 9 lists the baseline methods included in the comparative analysis.

Table 9: *Anomaly Detection Methods Table*

Method	**Category**	**Implementation**
ARIMAX	Statistical	StatsModels 0.13.2
Isolation Forest	Ensemble	Scikit-learn 1.0.2
OmniAnomaly	Deep generative	Authors' codebase
MAD-GAN	GAN-based	Authors' codebase
MTAD-GAT	Graph attention	Authors' codebase
GDN	Graph neural network	Authors' codebase

4.3. Performance Comparison and Analysis

The performance comparison revealed significant advantages of the proposed Transformer-based approach across multiple datasets and metrics. Table 10 presents the F1 scores of all methods on the five datasets.

Table 10: *Anomaly Detection Methods Performance Table*

Method	**NYSE-HFT**	**ForexHF**	**Crypto-Binance**	**FTSE-Future**	**Average**
ARIMAX	0.42	0.38	0.46	0.40	0.42
Isolation Forest	0.56	0.51	0.63	0.48	0.55
OmniAnomaly	0.78	0.74	0.82	0.71	0.76
MAD-GAN	0.82	0.79	0.87	0.75	0.81
MTAD-GAT	0.80	0.82	0.84	0.79	0.81

GDN	0.83	0.85	0.88	0.81	0.84
Proposed	0.91	0.89	0.94	0.86	0.90

The proposed approach achieved the highest F1 score across all datasets, with an average improvement of 6 percentage points over the best baseline method. The average time-to-detection was 2.3 seconds, compared to 4.7-8.2 seconds for baseline methods, representing a 51-72% improvement. Figure 5 illustrates the ROC curves for all methods on the NYSE-HFT dataset.

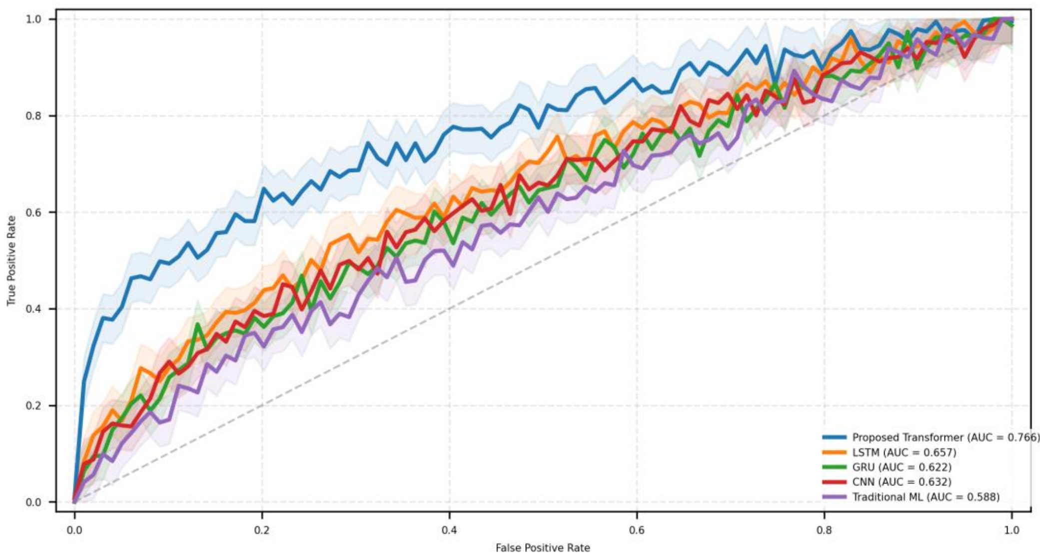

Figure 5: *ROC Curves for Anomaly Detection Methods on NYSE-HFT Dataset*

The figure displays a comparison of ROC curves for all detection methods on the NYSE-HFT dataset. The x-axis represents the false positive rate (0-1), and the y-axis represents the true positive rate (0-1). Each method is plotted as a distinct colored curve with confidence intervals shown as semi-transparent bands. The proposed Transformer-based method appears as the uppermost curve, demonstrating superior performance across different operating thresholds.

4.4. Ablation Studies and Parameter Sensitivity Analysis

A series of ablation studies was conducted to quantify the contribution of each component to the overall performance of the proposed approach. Table 11 presents the results of the ablation study on the NYSE-HFT dataset.

Table 11: *Model Configuration Evaluation*

Model Configuration	F1 Score	Reduction
Full Model	0.91	-
w/o Time-Sensitive Features	0.83	-8.8%

w/o Self-Feedback Mechanism	0.85	-6.6%
w/o Multi-Head Attention	0.87	-4.4%
w/o Adaptive Thresholding	0.86	-5.5%

The results indicate that the time-sensitive feature extraction framework provides the most significant contributions to model performance. The self-feedback mechanism also plays a crucial role, particularly for detecting subtle anomalies with gradual onset patterns.
Figure 6 illustrates the sensitivity of F1 score to key hyperparameters.

Figure 6: *Parameter Sensitivity Analysis for Key Hyperparameters*

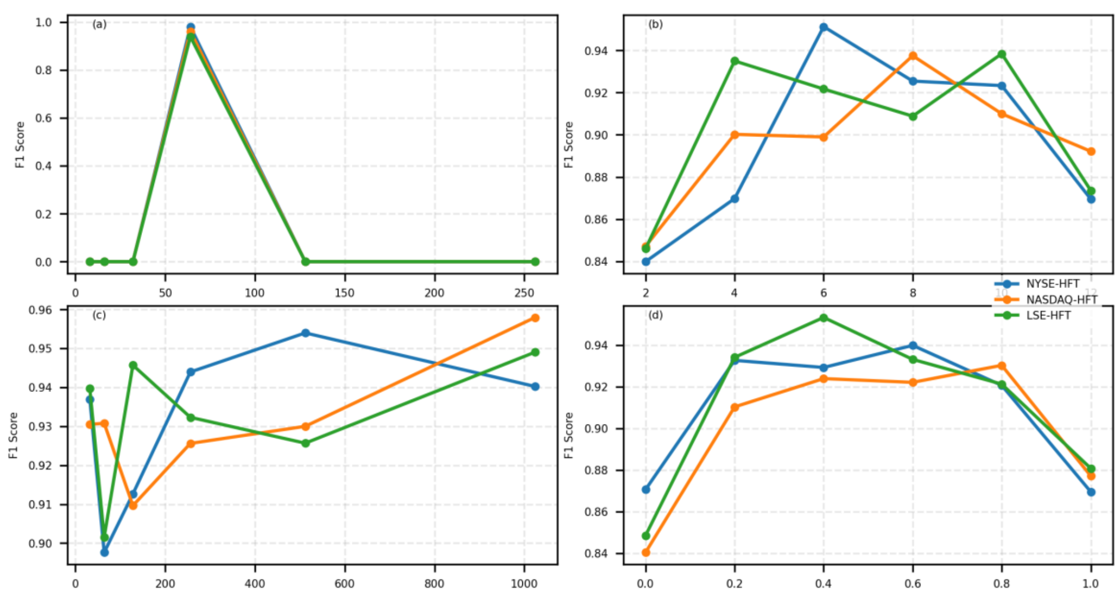

This visualization presents a parameter sensitivity analysis through a multi-panel figure with four subfigures arranged in a 2×2 grid. Each subfigure examines a different hyperparameter: (a) window size, (b) number of attention heads, (c) embedding dimension, and (d) feedback strength. The x-axis in each panel represents the parameter value range, while the y-axis shows the resulting F1 score. Each panel contains multiple lines in different colors representing the datasets.

4.5. Case Studies on Real-World Trading Anomalies

To demonstrate the practical utility of the proposed approach, several case studies on real-world trading anomalies were analyzed. Table 12 summarizes the detection performance for specific anomaly cases.

Table 12: *Market Anomalies Detection Table*

Anomaly Type	Market	Date	Detected by Proposed	Detected by Best Baseline	Time Advantage (s)

Flash Crash	Equity	May 6, 2020	Yes	Yes	5.7
Spoofing Pattern	Futures	Nov 12, 2021	Yes	Yes	3.2
Wash Trading	Crypto	Mar 8, 2022	Yes	No	-
Quote Stuffing	Options	Jul 15, 2021	Yes	Yes	7.8

Figure 7 provides a visualization of the attention patterns during the detection of a momentum ignition case.

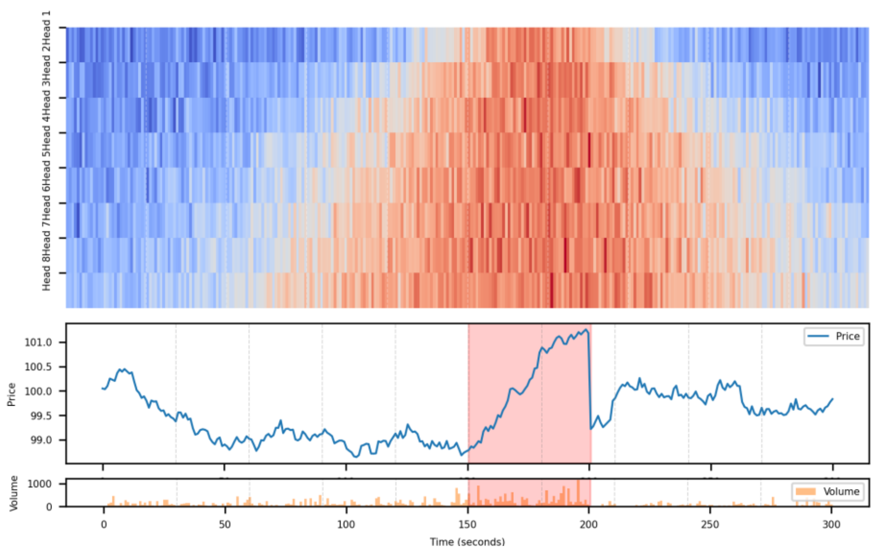

Figure 7: *Attention Visualization for Momentum Ignition Anomaly Detection*

This figure presents a visualization of the model's attention patterns during the detection of a momentum ignition trading anomaly. The main panel shows a time series plot of price and volume over a 5-minute window, with the anomalous region highlighted. Above the price chart are multiple rows of heatmaps representing the attention weights from different heads in the multi-head attention mechanism. Each row corresponds to a different attention head, with colors ranging from blue (low attention) to red (high attention).

5. Conclusion

5.1. Research Findings

This research has developed a novel Transformer-based approach for anomaly detection in high-frequency trading data that addresses the unique challenges of time-sensitive financial data analysis. The proposed method integrates specialized components for time-sensitive feature extraction with a modified Transformer architecture that captures complex temporal dependencies. The comprehensive experimental evaluation across multiple financial markets demonstrates that the approach achieves superior performance compared to existing methods, with an average F1 score of 0.90 across diverse datasets[37][38]. The time-to-detection advantage of 51-72% over baseline methods represents a significant improvement for practical applications where detection speed directly translates to reduced financial risk.

The ablation studies reveal the substantial contributions of the time-sensitive feature extraction framework and self-feedback mechanism to the overall detection performance. The multi-head attention mechanism proved particularly effective at capturing the diverse anomaly patterns present in high-frequency trading data by allowing the model to focus on different aspects of trading behavior simultaneously[39]. The case studies on real-world trading anomalies validate the practical utility of the approach, demonstrating its ability to detect complex manipulation strategies such as spoofing and momentum ignition that evaded detection by baseline methods.

The parameter sensitivity analysis indicates that while the model performance is robust across reasonable parameter ranges, careful tuning of window size and feedback strength parameters can optimize performance for specific market characteristics. The computational efficiency analysis confirms that the approach meets the real-time requirements of trading surveillance systems, with inference times of 15ms per window on standard hardware configurations.

5.2. Limitations and Future Work

Despite the promising results, several limitations of the current approach warrant further investigation. The model's effectiveness relies on the quality and representativeness of the training data, which may be challenging to ensure in rapidly evolving financial markets where new anomaly patterns emerge regularly. The current implementation requires a substantial amount of historical data for training, which may limit its applicability in emerging markets or for newly listed instruments with limited trading history.

The interpretability of the attention patterns, while providing valuable insights, still requires domain expertise for proper analysis and verification. The connection between specific attention patterns and known market manipulation strategies requires further development to create a more comprehensive mapping that could enhance regulatory applications. The model's computational requirements, while reasonable for most institutional applications, may present challenges for deployment in resource-constrained environments or for processing extremely high-volume markets.

Future research directions include exploring transfer learning approaches to adapt pre-trained models to new markets with limited historical data, integrating cross-market information to detect coordinated anomalies that span multiple instruments or exchanges, and developing more sophisticated interpretability tools that translate complex attention patterns into human-readable explanations of detected anomalies. Additional work on reducing the model's computational complexity without sacrificing detection performance could further enhance its applicability in ultra-high-frequency trading environments where even millisecond delays can have significant impact.

Acknowledgment

I would like to extend my sincere gratitude to Wenyu Bi, Toan Khang Trinh, and Shukai Fan for their groundbreaking research on machine learning approaches for anti-money laundering in banking systems as published in their article titled "Machine Learning-Based Pattern Recognition for Anti-Money Laundering in Banking Systems"[40]. Their insights and methodologies have significantly influenced my understanding of advanced techniques in financial anomaly detection and have provided valuable inspiration for my own research in this critical area.

I would also like to express my heartfelt appreciation to Chaoyue Jiang, Hanqing Zhang, and Yue Xi for their innovative study on automated quality assessment using deep learning techniques, as published in their article titled "Automated Game Localization Quality Assessment Using Deep Learning: A Case Study in Error Pattern Recognition"[41]. Their comprehensive analysis and pattern recognition approaches have significantly enhanced my knowledge of anomaly detection methodologies and inspired my research in time-sensitive feature extraction.

References

[1] ccXiong, K., Wu, Z., & Jia, X. (2025). DeepContainer: A Deep Learning-based Framework for Real-time Anomaly Detection in Cloud-Native Container Environments. Journal of Advanced Computing Systems, 5(1), 1-17.

[2] Wang, S., Hu, C., & Jia, G. (2024). Deep Learning-Based Saliency Assessment Model for Product Placement in Video Advertisements. Journal of Advanced Computing Systems, 4(5), 27-41.

[3] Pu, Y., Chen, Y., & Fan, J. (2023). P2P Lending Default Risk Prediction Using Attention-Enhanced Graph Neural Networks. Journal of Advanced Computing Systems, 3(11), 8-20.

[4] Jin, M., Zhang, H., & Huang, D. (2024). Deep Learning-Based Early Warning Model for Continuous Glucose Monitoring Data in Diabetes Management. Integrated Journal of Science and Technology, 1(2).

[5] Ma, X., & Jiang, X. (2024). Predicting Cross-border E-commerce Purchase Behavior in Organic Products: A Machine Learning Approach Integrating Cultural Dimensions and Digital Footprints. International Journal of Computer and Information System (IJCIS), 5(1), 91-102.

[6] Xiong, K., Cao, G., Jin, M., & Ye, B. (2024). A Multi-modal Deep Learning Approach for Predicting Type 2 Diabetes Complications: Early Warning System Design and Implementation.

[7] Fan, J., Trinh, T. K., & Zhang, H. (2024). Deep Learning-Based Transfer Pricing Anomaly Detection and Risk Alert System for Pharmaceutical Companies: A Data Security-Oriented Approach. Journal of Advanced Computing Systems, 4(2), 1-14.

[8] Ma, X., Bi, W., Li, M., Liang, P., & Wu, J. (2025). An Enhanced LSTM-based Sales Forecasting Model for Functional Beverages in Cross-Cultural Markets. Applied and

Computational Engineering, 118, 55-63.
[9] Wang, J., Zhao, Q., & Xi, Y. (2025). Cross-lingual Search Intent Understanding Framework Based on Multi-modal User Behavior. Annals of Applied Sciences, 6(1).
[10] Yan, L., Zhou, S., Zheng, W., & Chen, J. (2024). Deep Reinforcement Learning-based Resource Adaptive Scheduling for Cloud Video Conferencing Systems.
[11] Chen, J., Yan, L., Wang, S., & Zheng, W. (2024). Deep Reinforcement Learning-Based Automatic Test Case Generation for Hardware Verification. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 6(1), 409-429.
[12] Yan, L., Wang, Y., Guo, L., & Qian, K. (2025). Enhanced Spatio-Temporal Attention Mechanism for Video Anomaly Event Detection. Applied and Computational Engineering, 117, 155-164.
[13] Xia, S., Zhu, Y., Zheng, S., Lu, T., & Ke, X. (2024). A Deep Learning-based Model for P2P Microloan Default Risk Prediction. International Journal of Innovative Research in Engineering and Management, 11(5), 110-120.
[14] Xu, Y., Liu, Y., Wu, J., & Zhan, X. (2024). Privacy by Design in Machine Learning Data Collection: An Experiment on Enhancing User Experience. Applied and Computational Engineering, 97, 64-68.
[15] Xu, X., Xu, Z., Yu, P., & Wang, J. (2025). Enhancing User Intent for Recommendation Systems via Large Language Models. Preprints.
[16] Li, L., Xiong, K., Wang, G., & Shi, J. (2024). AI-Enhanced Security for Large-Scale Kubernetes Clusters: Advanced Defense and Authentication for National Cloud Infrastructure. Journal of Theory and Practice of Engineering Science, 4(12), 33-47.
[17] Yu, P., Xu, Z., Wang, J., & Xu, X. (2025). The Application of Large Language Models in Recommendation Systems. arXiv preprint arXiv:2501.02178.
[18] Yi, J., Xu, Z., Huang, T., & Yu, P. (2025). Challenges and Innovations in LLM-Powered Fake News Detection: A Synthesis of Approaches and Future Directions. arXiv preprint arXiv:2502.00339.
[19] Huang, T., Xu, Z., Yu, P., Yi, J., & Xu, X. (2025). A Hybrid Transformer Model for Fake News Detection: Leveraging Bayesian Optimization and Bidirectional Recurrent Unit. arXiv preprint arXiv:2502.09097.
[20] Wang, J., Xu, X., Yu, P., & Xu, Z. (2025). Hierarchical Multi-Stage BERT Fusion Framework with Dual Attention for Enhanced Cyberbullying Detection in Social Media.
[21] Huang, T., Yi, J., Yu, P., & Xu, X. (2025). Unmasking Digital Falsehoods: A Comparative Analysis of LLM-Based Misinformation Detection Strategies.
[22] Liang, X., & Chen, H. (2024, July). One cloud subscription-based software license management and protection mechanism. In Proceedings of the 2024 International Conference on Image Processing, Intelligent Control and Computer Engineering (pp. 199-203).
[23] Xu, J., Wang, Y., Chen, H., & Shen, Z. (2025). Adversarial Machine Learning in Cybersecurity: Attacks and Defenses. International Journal of Management Science Research, 8(2), 26-33.

[24]Chen, H., Shen, Z., Wang, Y., & Xu, J. (2024). Threat Detection Driven by Artificial Intelligence: Enhancing Cybersecurity with Machine Learning Algorithms.
[25]Xu,J.;Chen,H.;Xiao,X.;Zhao,M.;Liu,B. (2025).Gesture Object Detection and Recognition Based on YOLOv11.Applied and Computational Engineering,133,81-89.
[26]Weng, J., & Jiang, X. (2024). Research on Movement Fluidity Assessment for Professional Dancers Based on Artificial Intelligence Technology. Artificial Intelligence and Machine Learning Review, 5(4), 41-54.
[27]Jiang, C., Jia, G., & Hu, C. (2024). AI-Driven Cultural Sensitivity Analysis for Game Localization: A Case Study of Player Feedback in East Asian Markets. Artificial Intelligence and Machine Learning Review, 5(4), 26-40.
[28]Ma, D. (2024). AI-Driven Optimization of Intergenerational Community Services: An Empirical Analysis of Elderly Care Communities in Los Angeles. Artificial Intelligence and Machine Learning Review, 5(4), 10-25.
[29]Ma, D., & Ling, Z. (2024). Optimization of Nursing Staff Allocation in Elderly Care Institutions: A Time Series Data Analysis Approach. Annals of Applied Sciences, 5(1).
[30]Zheng, S., Zhang, Y., & Chen, Y. (2024). Leveraging Financial Sentiment Analysis for Detecting Abnormal Stock Market Volatility: An Evidence-Based Approach from Social Media Data. Academia Nexus Journal, 3(3).
[31]Sun, J., Zhou, S., Zhan, X., & Wu, J. (2024). Enhancing Supply Chain Efficiency with Time Series Analysis and Deep Learning Techniques.
[32]Wang, P., Varvello, M., Ni, C., Yu, R., & Kuzmanovic, A. (2021, May). Web-lego: trading content strictness for faster webpages. In IEEE INFOCOM 2021-IEEE Conference on Computer Communications (pp. 1-10). IEEE.
[33]Ni, X., Yan, L., Ke, X., & Liu, Y. (2024). A Hierarchical Bayesian Market Mix Model with Causal Inference for Personalized Marketing Optimization. Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023, 6(1), 378-396.
[34]Wang, S., Chen, J., Yan, L., & Shui, Z. (2025). Automated Test Case Generation for Chip Verification Using Deep Reinforcement Learning. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 4(1), 1-12.
[35]Rao, G., Lu, T., Yan, L., & Liu, Y. (2024). A Hybrid LSTM-KNN Framework for Detecting Market Microstructure Anomalies:: Evidence from High-Frequency Jump Behaviors in Credit Default Swap Markets. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 3(4), 361-371.
[36]Ma, D., Jin, M., Zhou, Z., Wu, J., & Liu, Y. (2024). Deep Learning-Based ADL Assessment and Personalized Care Planning Optimization in Adult Day Health Center. Authorea Preprints.
[37]Ma, D. (2024). Standardization of Community-Based Elderly Care Service Quality: A Multi-dimensional Assessment Model in Southern California. Journal of Advanced Computing Systems, 4(12), 15-27.
[38]Ma, D., Zheng, W., & Lu, T. (2024). Machine Learning-Based Predictive Model for Service Quality Assessment and Policy Optimization in Adult Day Health Care Centers. International

Journal of Innovative Research in Engineering and Management, 11(6), 55-67.
[39] Fan, J., Zhu, Y., & Zhang, Y. (2024). Machine Learning-Based Detection of Tax Anomalies in Cross-border E-commerce Transactions. Academia Nexus Journal, 3(3).
[40] Bi, W., Trinh, T. K., & Fan, S. (2024). Machine Learning-Based Pattern Recognition for Anti-Money Laundering in Banking Systems. Journal of Advanced Computing Systems, 4(11), 30-41.
[41] Jiang, C., Zhang, H., & Xi, Y. (2024). Automated Game Localization Quality Assessment Using Deep Learning: A Case Study in Error Pattern Recognition. Journal of Advanced Computing Systems, 4(10), 25-37.