

Review

Not peer-reviewed version

Critical Analysis on Anomaly Detection in High-Frequency Financial Data Using Deep Learning for Options

[Sunil Mahadev Sangve](#)^{*} and [Soham Vivek Borkar](#)^{*}

Posted Date: 5 May 2025

doi: 10.20944/preprints202505.0168.v1

Keywords: machine learning; deep learning; recurrent neural network; graph neural network; anomaly detection; high frequency trading



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Review

Critical Analysis on Anomaly Detection in High-Frequency Financial Data Using Deep Learning for Options

Sunil Sangve and Soham Borkar

Dept. of Artificial Intelligence and Data Science, Vishwakarma Institute of Technology, Pune, India

* Correspondence: sunil.sangave@vit.edu

Abstract: High-frequency financial markets churn out massive amounts of data filled with intricate microstructural patterns, which makes them vulnerable to issues like spoofing, layering, and market manipulation. Traditional methods for detecting anomalies often fall short when it comes to capturing these complex patterns, especially given the fast-paced and ever-changing nature of financial transactions. In this study, we introduce a deep learning-based framework designed for spotting anomalies in high-frequency trading (HFT) data. This framework utilizes advanced techniques such as graph neural networks (GNN), recurrent neural networks (RNN), and transformer-based autoencoders. By integrating multi-modal feature extraction, attention to temporal dependencies, and adaptive learning strategies, we aim to boost detection accuracy. We tested our approach on real-world high-frequency limit order book (LOB) data, along with synthetic anomalies added to the dataset. The results from our experiments show notable enhancements in anomaly detection accuracy, precision, and recall, surpassing existing methods while keeping false positive rates low. Our findings underscore the promise of deep learning models in enhancing market surveillance, ensuring regulatory compliance, and mitigating financial risks in HFT settings.

Keywords: high-frequency trading (HFT); anomaly detection; deep learning; graph neural networks (GNN); recurrent neural networks (RNN); transformer autoencoder; financial market manipulation; limit order book (LOB); machine learning; time-series analysis

I. Introduction

The rapid advancements in financial markets, especially in high-frequency trading (HFT), have brought significant challenges to anomaly detection. Financial markets generate vast volumes of transaction data at sub-second intervals, increasing the difficulty of identifying manipulative activities and ensuring market integrity. Recent developments in deep learning and machine learning methodologies have offered promising solutions for anomaly detection in this complex environment. Numerous studies have explored diverse approaches, from empirical modeling to deep learning frameworks, demonstrating the evolving nature of anomaly detection techniques.

Brugman [1] developed a real-time monitoring system for fatigue detection in truckers, highlighting the significance of real-time anomaly detection in high-velocity environments. Although the focus was on transportation, the principles of real-time monitoring, anomaly detection, and decision support systems have strong relevance to high-frequency financial trading environments. Similarly, the ability to act on anomalies as they occur is critical for mitigating risks and ensuring system reliability.

Veryzhenko et al. [2] explored spoofing detection in HFT by leveraging machine learning techniques, particularly k-nearest neighbors (KNN). Their study demonstrated that machine learning models can effectively detect intraday manipulative behaviors in microsecond-scale trading data. However, it also revealed critical limitations such as sensitivity to training data nuances and

challenges posed by non-stationarity in financial time series. The importance of using supervised and unsupervised models for financial fraud detection emerged as a recurring theme, emphasizing the need for robustness and adaptability in anomaly detection frameworks.

Rizvi et al. [3] proposed an unsupervised approach using Empirical Mode Decomposition (EMD) combined with Kernel Density Estimation (KDE) clustering to detect stock price manipulation. Their methodology, which involved artificially injected anomalies into real datasets, achieved a maximum Area Under Curve (AUC) improvement of 84%. This work highlighted the strengths of unsupervised learning in dealing with unlabeled and highly volatile financial datasets, showcasing the potential of empirical techniques when labeled data is scarce or unreliable.

In another significant contribution, Zhao et al. [4] introduced a graph attention network (GAT)-based framework for multivariate time-series anomaly detection. Their model, MTAD-GAT, effectively captured complex inter-feature dependencies and temporal patterns across datasets such as SMAP, MSL, and TSA. Deep learning models such as GATs based on graph neural network were proven to enhance anomaly detection capacity over traditional sequential models to enable improved diagnosis and explainability. The developments represented a paradigm shift where the GNNs can represent interconnected streams of financial information more effectively.

Mejri et al. [5] conducted a comprehensive evaluation of unsupervised anomaly detection methods for time-series data, emphasizing the need for more comprehensive evaluation frameworks. Their research went beyond typical metrics by proposing specialized evaluation metrics for time-series anomalies, such as model stability and anomaly type differentiation. These types of evaluations are especially critical in finance, where anomaly types may vary greatly under varying market conditions and datasets.

With multimodal dataset fusion as the basis, Chen et al. [6] proposed a multi-modal graph neural network-based market manipulation detection system. By merging heterogeneous sources like price movements and order book dynamics with state-of-the-art attention mechanisms, the model achieved an astonishing 98.7% accuracy. The research demonstrated the worth of combining heterogeneous data modalities to represent complex and often concealed manipulation patterns in high-frequency finance data, paving the way for highly responsive and precise fraud detection systems.

Xi et al. [7] extended the application of deep learning to cross-border logistics further by integrating real-time multimodal data and anomaly detection with deep reinforcement learning (DRL). In the context of logistics, the principles of real-time decision-making, anomaly discovery, and adaptive optimization can be directly applied to financial trading systems where timely and responsive response to anomalous activity is the key issue.

Finally, Guanghe et al. [8] proposed innovative transformer networks for real-time dark pool trading anomaly detection. The system achieved an impressive 97.8% detection rate and more than 45% processing latency reduction. Utilization of sophisticated attention mechanisms and streamlined processing pipelines in transformer models demonstrated them to be more effective in coping with high-frequency, low-latency financial data environments, setting new standards of performance for market surveillance technology.

Together, the papers reflect the changing dynamic of anomaly detection approaches in high-frequency finance environments. They highlight the movement away from traditional statistical methods toward more capable deep models for handling the temporal, structural, and multimodal complexities of financial data. As outstanding as the advancements have been, model interpretability, resilience against novel patterns of fraud, and efficiency at scale continue to be dominant drivers of research in this crucial area of financial data analysis.

II. Literature Review

The high-frequency trading (HFT) boom has led to the need for sophisticated anomaly detection approaches to counter manipulative strategies in the markets and provide financial integrity. The following critically examines recent work on deep learning-based approaches, such as methods,

datasets, evaluation metrics, and their limitations. The papers considered indicate that neural networks, reinforcement learning, and graph models are capable of identifying anomalies with limitation in interpretability, scalability, and real-time capability.

Cao et al. [9] proposed a new paradigm based on deep reinforcement learning (DRL) to optimize HFT strategies. They combined convolutional neural networks with short-term memory models to capture both short-term and long-term market trends. Trained on tick-level NASDAQ data, their model achieved a Sharpe ratio of 3.42, outperforming traditional benchmarks by 33%. The DRL agent exhibited resilience across volatile and stable market periods, mirroring its capacity to generalize. Advanced data preprocessing was critical in enhancing model learning efficiency. However, the study noted constraints like excessive computational requirements for resources and low generalizability to other markets or time horizons. Furthermore, while financial performance indicators were maximized, other key considerations like regulatory compliance and model interpretability were not prioritized. Nevertheless, the study set precedent for further studies with a goal to integrate DRL within dynamic and complex financial market conditions [9].

Li et al. [10] addressed the challenges of anomaly pattern detection in high-frequency trading by proposing a graph neural network (GNN)-based method. Using the LOBSTER dataset, which includes top stocks like Apple and Amazon, their model efficiently extracted intricate trading patterns using graph-based features and attention mechanisms. Experimental results revealed a 15% improvement in detection accuracy over traditional models, in addition to real-time processing efficiency essential for modern surveillance systems. The architecture itself possessed good robustness towards market volatility and showed its ability to detect both local and global anomalies. However, the study admitted that there were restrictions regarding the handling of high-dimensional data and dynamic mutation of fraud schemes, which could affect long-term performance. The method's focus on low computational overhead and capacity to identify complex trading patterns constituted great advancements in financial anomaly detection through deep learning methodologies [10].

Alaminos et al. [11] proposed a hybrid model combining ARMA-GARCH econometric techniques and deep neural networks to examine intraday trading tactics, particularly amid geopolitical turmoil. Focusing on defense sector equities and FOREX fluctuations, the hybrid model recorded high levels of prediction, where ARMA-GARCH-QRNN stood out at 93.59%. Their work identified the resistance of hybrid models to unstable market conditions by accurately capturing nonlinear and nonstationarity trends. Interestingly, the model identified potential investment strategies in periods of global uncertainty. Despite such successes, the strategy was faced with computational challenges, particularly due to the complexity of quantum and deep neural elements. Additionally, sudden market shocks posed difficulties in generating accurate predictions using the ARMA-GARCH model. Nonetheless, this research confirmed the advantage of hybrid methods for enhancing forecasting performance in high-frequency, high-volatility conditions relevant to financial anomaly detection and decision-making models [11].

Shanmuganathan and Suresh [12] designed a Markov-augmented Long Short-Term Memory (I-LSTM) approach to enhance anomaly detection in sequential time-series data. Their solution, tested on industrial sensor readings like temperature and humidity, indicated smaller prediction errors compared to the standard algorithms of KNN and ordinary LSTM models. Markov modeling complemented with sophisticated LSTM structures further improved sequential learning by making the algorithm suitable for latent anomaly detection in streaming data environments. While their intended field was industrial monitoring, the method yielded valuable conclusions applicable to financial markets with continuous, high-speed, and intricate data streams. The authors stressed challenges emanating from a lack of pre-labeled anomaly data and underscored the importance of strong preprocessing. Additionally, they identified that although predictive accuracy was improved, scalability and handling of various real-world cases remained future work, particularly for high-frequency and high-noise situations such as trading [12].

Bello et al. [13] provided an extensive review on the integration of deep learning models in real-time fraud detection for high-frequency trading (HFT) environments. Their analysis focused on CNNs and RNNs, advocating for hybrid approaches that combine deep learning capabilities with traditional statistical models. The study stressed the critical role of model retraining, explainable AI, and robust preprocessing in maintaining fraud detection efficacy amidst evolving market conditions. Challenges such as handling vast data volumes, ensuring computational efficiency, and maintaining low false positive rates were emphasized. The authors also pointed out that model transparency is necessary for regulatory compliance and user trust. Future research directions included the adoption of ensemble models, continuous online learning frameworks, and advanced data augmentation strategies. By synthesizing conceptual and technical challenges, this paper underscored the importance of resilient, adaptive deep learning architectures in the increasingly complex and fast-paced environment of HFT fraud detection [13].

Li et al. [14] investigated network traffic anomaly detection in IoT environments, presenting a deep learning framework that achieved remarkable performance, including 99.2% accuracy and 0.998 AUC-ROC on the IoT-23 dataset. While primarily targeting IoT systems, the model—utilizing a CNN-LSTM hybrid and stacked autoencoders—provided transferable insights into handling large-scale, high-frequency datasets with heterogeneous structures. The study emphasized the importance of combining manual feature engineering with deep unsupervised learning for robust anomaly detection. However, challenges included high computational costs, the necessity for continual retraining, and difficulties in maintaining model interpretability—issues that are directly relevant to financial data analysis. Notably, the study stressed that zero-day attack detection and adaptability to novel patterns remain critical research areas. The focus on Grad-CAM-based interpretability methods was also crucial, suggesting pathways for increasing transparency in anomaly detection models, an essential aspect for regulatory applications in financial domains [14].

Jin et al. [15] presented a broad review on the use of Graph Neural Networks (GNNs) in applications including time-series forecasting, classification, imputation, and anomaly detection, referring to it as GNN4TS. Jin et al. classified existing work by task-related objectives and structural methodology, showcasing structured coverage for GNN-driven techniques. The survey highlighted advantages in modeling spatial-temporal relationship using graph-based methods and subsequently bridging weaknesses from conventional sequential models. Yet, the research pinpointed issues such as architecture complexity, choosing hyperparameters, and failing to quantify uncertainty in current GNN models. The researchers suggested utilizing AutoML-guided design in order to automate model construction workflows. Mentioned as well was the need for expanding the GNN to non-Euclidean and noisy data, the detection of a fundamental gap in real-world adoption, particularly with regards to high-frequency finance data. This study provided valuable insights to researchers who wanted to apply GNNs in dynamic and complicated settings, like financial anomaly detection [15].

Haq et al. [16] introduced HF-PPAD, an automated machine learning framework for detecting anomalous peak patterns in hydrological time-series data. By generating synthetic anomalies using TimeGAN and applying models like Temporal Convolutional Networks (TCN), InceptionTime, and LSTM, the system significantly enhanced anomaly detection without requiring deep domain knowledge. Although the primary application was environmental monitoring, the focus of the approach on automatic model selection, hyperparameter optimization, and synthetic data generation has direct applications to financial markets. The main issues were reliance on synthetic anomalies, generalization to noisy real-world data being poor, and computational cost at training time. Yet, HF-PPAD's capability to improve anomaly detection accuracy without human involvement is a significant achievement for high-frequency financial anomaly detection, where model deployment on time and adaptation to evolving conditions are essential [16].

Bakumenko and Elragal [17] explored applying machine learning algorithms to identify anomalies in financial data with the purpose of improving fraud detection accuracy. Using a highly unbalanced set of normal and anomalous transactions, they applied supervised techniques like

Random Forest and unsupervised techniques like Isolation Forest with notable improvements in recall rates. They called out for better preprocessing quality, i.e., journal entry vectorization and hyperparameter tuning. Constraints identified included class imbalance, challenging to accommodate feature variability, and long journal entry issues. Specifically, the paper stressed avoiding false negatives for fraud detection in finance. This work reaffirmed the necessity of robust machine learning pipelines for anomaly detection in high-frequency financial systems, particularly with financial data complexity and volumes continuing to increase [17].

Poutré et al. [18] proposed a deep unsupervised anomaly detection system for high-frequency markets based on a transformed Transformer autoencoder architecture. Their approach learned subsequences of the limit order book (LOB) over time to model normal behavior and employed a dissimilarity function for anomaly detection for out-of-sample data. Synthetic manipulations such as quote stuffing and pump-and-dump attacks were incorporated to test model robustness. The model performed improved fraud detection, outperforming traditional unsupervised methods in precision and F1-score. The authors, however, acknowledged limitations such as interpretability limitations, modeling collective anomalies difficulties, and computational expense under extreme market conditions. The work also highlighted the need for dynamic adaptation to evolving fraud tactics, which remains a top challenge in real-time financial systems. Overall, this work provided a good guideline for constructing scalable, unsupervised anomaly detection models in high-frequency trading environments [18].

Yang et al. [19] proposed a federated learning (FL) framework to enable privacy-preserving sharing of medical data for use in drug development. Their approach utilized decentralized machine learning techniques to allow multiple institutions to collaboratively train models without seeing raw patient data, thereby maintaining confidentiality and complying with data protection regulations. The study introduced enhanced aggregation methods to mitigate data heterogeneity issues across different medical centers, resulting in significantly improved model accuracy and robustness. Although focused on healthcare, the federated learning privacy-preserving principles are very translatable to finance applications, where data sensitivity and confidentiality are paramount. Communication overhead, convergence instability under the existence of non-independent and identically distributed (non-IID) data, and vulnerability to adversarial attacks were main issues that were identified. Nevertheless, the research demonstrated that FL could be a robust foundation for secure, distributed anomaly detection systems in applications like high-frequency trading, where data privacy among multiple entities and facilitating collaborative learning are becoming increasingly vital [19].

Guanghe et al. [20] designed a real-time anomaly detection model for dark pool trading on an enhanced Transformer network. By integrating maximized attention mechanisms with cutting-edge streaming analytics, their system achieved 97.8% accuracy of detection using merely a 2.3-millisecond latency of processing. Emphasizing dark pools that are typically black and vulnerable to manipulation, the system achieved unprecedented performance levels in market surveillance solutions. Nonetheless, limitations such as intensive computational requirements for resources as well as their dependency on prelabeled history for training were cited by authors. Flexibility to evolving trends in fraud and low latency amidst high market volatility were among the listed future issues. Despite these, the research demonstrated how more advanced Transformer-based models had potential to make considerable improvements in anomaly detection in highly sensitive, low-visibility trading scenarios, highlighting the necessity to adopt real-time, deep learning-based solutions for financial market integrity [20].

Ashtiani and Raahemi [21] conducted a comprehensive systematic literature review focused on intelligent fraud detection in financial statements using machine learning (ML) and data mining techniques. By analyzing over 100 peer-reviewed studies, the authors categorized methods into supervised, unsupervised, and hybrid approaches, noting that ensemble learning models like Random Forest and boosting algorithms consistently achieved superior performance in fraud detection tasks. The review highlighted key challenges such as severe class imbalance, feature

selection complexity, lack of publicly available fraud datasets, and the dynamic nature of fraudulent behavior. They emphasized the growing importance of explainability and transparency in ML models, especially under increasing regulatory scrutiny. Although primarily centered on corporate financial fraud, the study's insights into data preprocessing, model robustness, and performance evaluation are directly applicable to anomaly detection in high-frequency financial data environments. Their findings suggest that hybrid models combining financial knowledge and machine learning yield the most reliable fraud detection systems [21].

Wang et al. [22] explored the synergy of deep learning and cloud computing to improve personal search, designing a hybrid architecture that registered impressive MAP improvement and query response times. Their system reported a 15% MAP gain while guaranteeing sub-200ms latency for 95% of the queries, emphasizing its scalability and efficiency. The research used deep neural networks combined with hierarchical monitoring networks to provide adaptive personalization according to user behavior. Although highlighting benefits of AI-powered personalization, the authors also brought up concerns about privacy, data reliance, and computational expense of using deep models at scale. The research indicated possible "filter bubble" effects where the users may be confined to known information streams, thereby reducing search diversity. Though constrained by these limits, the study showed significant advancements in real-time user-focused search applications and presented useful lessons for constructing scalable AI infrastructures—principles that apply to real-time financial anomaly detection systems [22].

Wang et al. [23] studied distributed high-performance computing (HPC) methods for speeding up deep learning training, with architectures such as model parallelism and data parallelism. Their study showed that hybrid approaches integrating decentralized communication with adaptive gradient optimization greatly improved scalability and efficiency of deep learning models. Methods such as Ring All Reduce and 1-bit stochastic gradient descent were highlighted for reducing communication bottlenecks, essential for real-time model updates in large-scale systems. The authors also pointed out challenges such as fault tolerance, energy efficiency, and load balancing in distributed training. While mainly aimed at deep learning applications, the observations of the paper directly translate to ultra-fast retraining requirements in real-time financial systems that need to quickly retrain models in order to fit changing environments. Moreover, the study highlighted that breakthroughs in GPU clusters and dedicated AI accelerators could pave the way for the next wave of low-latency, high-precision anomaly detection systems, thereby giving great relevance to financial market applications needing ultra-fast processing [23].

Ibikunle et al. [24] examined the use of machine learning (ML) to reveal novel insight into high-frequency trading (HFT) practices. They utilized tree-based ensembles and a multi-target method to examine liquidity-supplying and liquidity-demanding behaviors in HFT. Based on NASDAQ-provided data, they revealed that ML models were successful at capturing nonlinear relationships and were better at explaining the effect of HFT on price discovery. Most notably, liquidity-supplying strategies were shown to react more sensitively to information events compared to liquidity-demanding ones. Nevertheless, the paper also identified major limitations, such as a lack of granularity in existing HFT datasets and possible information biases based on general trading metrics. The authors stressed more precise data to actually distinguish between different HFT strategies. In general, this work highlighted the potential of ML to revolutionize the comprehension and optimization of high-frequency market behavior, providing avenues for improved anomaly detection and improved market regulation via machine-learning-based analytics [24].

Zhang and Hua [25] introduced a detailed survey discussing key challenges in the analysis of high-frequency financial data, such as nonstationarity, low signal-to-noise ratios, and asynchronous data. Analyzing more than 150 articles, they classified solutions into data preprocessing techniques (such as differencing and smoothing) and quantitative modeling approaches (such as adaptive algorithms and deep neural networks). Their research specifically highlighted that although new algorithms hold potential, threats such as overfitting, poor generalization, and higher computational requirements are ongoing. Notably, they supported methods that cater to intraday seasonality and

enhance synchronization within transaction records for better forecasting accuracy. The article's elaborate charts and taxonomies offered a useful guide for researchers addressing HFT challenges. Notwithstanding the improvements described, they acknowledged that most advanced techniques falter with real-world out-of-sample data. Their observations reiterated the importance of strong preprocessing, adaptive modeling, and ongoing validation when applying anomaly detection methods in high-frequency financial settings [25].

Bello et al. [26] investigated adaptive machine learning models for real-time fraud detection in dynamic financial environments, focusing on ongoing learning from shifting transaction patterns. Methods such as reinforcement learning, online learning, and explainable AI (e.g., SHAP, LIME) were incorporated into their fraud detection architectures to enable model transparency and adaptability. Their work highlighted the importance of high-quality transactional data, real-time analytics infrastructures (e.g., Apache Kafka and Flink), and computational efficiency for successful fraud prevention. They also touched on ethical issues, noting the importance of ensuring fairness and reducing bias in financial systems. Issues like imbalanced datasets, the need for low-latency responses, and computational overhead were identified as persistent research challenges. In general, the paper emphasized the need for robust, scalable, and explainable machine learning models that can adapt to new trends in fraud over time, an objective directly transferrable to improving anomaly detection functionality in high-frequency, ever-changing financial markets [26].

Basit et al. [27] studied the use of Quantum Variational Autoencoders (QVAEs) to improve predictive analytics in high-frequency trading (HFT) markets with the view to improving market anomaly detection. Through the use of quantum computing, their research proved that QVAEs performed better than traditional models in accuracy, recall, and F1-score. Based on Dow Jones and NASDAQ data, the quantum models exhibited greater precision-recall AUC values, which reflected better anomaly detection performance in high-dimensional, complex trading data. However, the research admitted significant limitations, such as training instability, longer optimization times, and existing hardware limitations in quantum computing. Furthermore, fine-tuning quantum models proved challenging, requiring specialized techniques to stabilize learning processes. Despite these hurdles, the paper provided critical evidence supporting the potential of quantum-enhanced models for future HFT applications, especially in the realm of detecting sophisticated and evolving market anomalies that traditional methods often fail to capture [27].

Nguyen et al. [28] proposed two deep learning-based approaches—an LSTM network for forecasting and an LSTM autoencoder combined with one-class SVM for anomaly detection—specifically targeting supply chain management applications. By utilizing both external and internal data sources, their models achieved superior performance compared to traditional methods, particularly in predicting remaining useful life (RUL) and detecting anomalous sales patterns. The study validated its methods using C-MAPSS datasets and fashion retail sales data. Although the application domain was supply chains, the dual approach of forecasting and anomaly detection had direct parallels to financial systems, where predictive analytics and anomaly identification are crucial for risk mitigation. Key limitations included the computational demands of LSTM-based models and challenges in hyperparameter optimization. Nonetheless, the paper highlighted the strong potential of hybrid deep learning frameworks for real-world time-series anomaly detection, contributing valuable methodologies translatable to high-frequency financial data analysis and fraud prevention [28].

Zhao et al. [29] introduced DeepLOB, a deep convolutional neural network (CNN) framework designed for modeling and predicting limit order book (LOB) data. The model captures both spatial and temporal dependencies inherent in high-frequency trading data, outperforming traditional statistical models and shallow machine learning algorithms. DeepLOB leverages CNN layers to extract hierarchical representations from raw LOB inputs without manual feature engineering. Their results showed significant improvements in mid-price movement prediction, with higher accuracy and reduced latency, critical factors for real-time trading applications. Although promising, the authors noted that DeepLOB's effectiveness is heavily dependent on the quality and granularity of

input data, and challenges persist in generalizing the model across various asset classes and trading environments. Nevertheless, this study represents a major advancement in applying deep learning architectures to financial market microstructure analysis and high-frequency anomaly detection scenarios [29].

Sirignano and Cont [30] applied deep learning techniques, specifically recurrent neural networks (RNNs), to study price formation in large-scale financial markets using high-frequency trading data. By training models on billions of order book events, they uncovered universal patterns underlying asset price movements, demonstrating that deep learning could detect complex dynamics that traditional econometric methods failed to capture. Their findings show that some of the underlying structures in high-frequency trading data are common across markets and assets, further pointing towards the potential for transfer learning. Scalability, interpretability, and generalizability across market regimes were identified as ongoing challenges. This research greatly emphasized that deep learning is not only a predictive technique but can also be employed to detect underlying mechanisms in financial markets, thereby enhancing the efficiency of anomaly detection systems for high-frequency environments [30].

Zhang et al. [31] proposed Temporal Graph Attention Networks (TGAT) for spatio-temporal modeling of financial transaction data. TGAT combines graph neural networks (GNNs) with temporal attention mechanisms for capturing the relational dependencies and dynamic patterns of high-frequency trading situations simultaneously. Their proposed model had superior predictive performance compared to standard time-series and static graph models for fraud detection and anomaly prediction tasks. The paper also highlighted the key significance of flexible models addressing dynamic data distribution and structural change in trading networks. The contribution to note was the presentation of the fact that dynamic graph usage considerably improves anomaly detection scores compared to conventional fixed forms. The paper nonetheless addressed limitations that encompass computationally scalability, interpretability. However, on the whole, TGAT offers a better architecture in representation of rich evolving anomalies present in high-frequency finance [31].

Zhu and Chan [32] explored applying Long Short-Term Memory (LSTM) networks to identify anomalies in high-frequency trading data. In their research, the authors demonstrated that LSTMs could efficiently extract temporal relationships and recognize anomalous activities such as spoofing and layering in limit order books. The authors employed a combination of actual trading data and synthetically added anomalies for training and testing the LSTM models. Outcomes included higher detection rates and lower false positives compared to traditional threshold-based and rule-based mechanisms. Nevertheless, they added that the model proved to be sensitive to hyperparameters and performance equally sensitive to how the training set was balanced. The article put a focus on the importance of the necessity for effective dataset construction and the potential of recurrent deep learning models to significantly enhance market surveillance and fraud detection in HFT settings. The findings directly contribute to the growing field of deep learning applications in financial anomaly detection [32].

Comparative analysis of existing studies on anomaly detection in high-frequency financial data are summarized in Table 1.

Table 1. Comparative Analysis of Existing Studies on Anomaly Detection in High-Frequency Financial Data.

Sr . N o.	Paper (First Author + Year)	Problem Addressed	Methods Used	Dataset Used	Key Findings	Limitations
-----------	-----------------------------	-------------------	--------------	--------------	--------------	-------------

1	Brugman (2022)	Real-time fatigue detection in trucking industry	Real-time monitoring system, CNN	Physiological data from truckers	Demonstrated feasibility of real-time fatigue monitoring	Limited to physiological signals, not financial domain
2	Veryzhenko (2023)	Spoofing detection in high-frequency trading (HFT)	Machine Learning classifiers (SVM, Random Forest)	Simulated order book data	Improved spoofing detection performance with ML models	Absence of real HFT trading data for validation
3	Rizvi (2018)	Stock price manipulation detection	Empirical Mode Decomposition (EMD) + Kernel Density Estimation (KDE) Clustering	Simulated financial time series	Proposed an effective clustering approach for unsupervised manipulation detection	Scalability and real-time deployment not addressed
4	Zhao (2020)	Multivariate time-series anomaly detection	Graph Attention Networks (GAT)	Public industrial datasets (not specific to finance)	Outperformed traditional time-series models using graph-based approach	Model complexity; real-time inference not tested
5	Mejri (2024)	Evaluation of time-series anomaly detection methods	Comparative study of 10+ algorithms (Forecasting, Reconstruction)	Public time-series datasets	Provided detailed benchmarking of unsupervised techniques	Lack of focus on financial or high-frequency datasets
6	Chen (2024)	Market manipulation detection using multi-modal data	Graph Neural Networks (GNN) on order book and trade network data	Simulated multi-modal HFT datasets	Achieved better detection by combining multiple data modalities	Scalability to live data streams remains a challenge
7	Xi (2024)	Route optimization and anomaly detection in logistics	Deep Reinforcement Learning (DRL)	Cross-border logistics datasets	Proposed a multimodal DRL framework	Application outside finance; financial anomalies not considered
8	Guanghe (2024)	Real-time anomaly detection in dark pool trading	Enhanced Transformer Networks	Simulated dark pool transaction data	Improved performance for opaque trading environments	Lack of real dark pool datasets for training and validation

9	Cao (2024)	Optimization of HFT strategies	Deep Reinforcement Learning (DRL)	Simulated HFT trading datasets	DRL models achieved superior performance in dynamic strategy optimization	Focused on strategy optimization, not anomaly detection
10	Li (2024)	Anomaly pattern detection in HFT	Graph Neural Networks	Financial transaction datasets	Successfully captured complex transaction patterns for anomaly detection	Scalability and real-time deployment issues
11	Alaminos (2024)	Intraday strategy exploration in HFT	Hybrid ARMA-GARCH-Neural Network model	High-frequency trading data	Demonstrated improved intraday trading strategy prediction using hybrid models	Limited real-time adaptability; offline model focus
12	Shanmuganathan (2024)	Anomaly detection in time-series sensor data	Markov-enhanced I-LSTM	Sensor data (non-financial)	Enhanced anomaly detection accuracy through Markov modeling	Not directly validated on financial datasets
13	Bello (2024)	Fraud detection in high-frequency trading	Deep Learning architectures	Conceptual framework, no specific dataset used	Addressed conceptual challenges and proposed real-time fraud detection solutions	Lack of empirical results and benchmarks
14	Li (2024)	Network traffic anomaly detection in IoT environments	Deep Learning models (CNN, LSTM)	IoT network datasets	Showed effectiveness of DL models for network anomaly detection	Application domain outside of financial trading
15	Jin (2024)	Survey of GNNs for time-series tasks	Review of forecasting, classification, imputation, anomaly detection using GNNs	Various time-series datasets	Summarized the potential of GNNs across multiple time-series applications	Limited specific financial use-case demonstrations
16	Haq (2024)	Detecting anomalous peak patterns in watershed time series	Automated ML (AutoML) approach	Watershed environmental datasets	Demonstrated effectiveness of AutoML for peak anomaly detection	Application domain is environmental science, not finance

17	Bakume nko (2022)	Financial data anomaly detection using ML	Machine Learning algorithms (XGBoost, RF, SVM)	Financial transaction datasets	ML models successfully identified financial anomalies	Need for higher-dimensional and real-time adaptation
18	Poutré (2024)	Deep unsupervised anomaly detection in HFT markets	Deep Autoencoder models	High-frequency market data	Achieved strong anomaly detection without labeled data	Challenges in scaling for large real-time systems
19	Yang (2024)	Privacy-preserving data sharing using federated learning	Federated Learning frameworks	Medical datasets (non-financial)	Demonstrated privacy-preserving anomaly detection using federated learning	No application to financial markets or HFT scenarios
20	Cao (2024)	Real-time anomaly detection in dark pool trading (repeated entry)	Enhanced Transformer Networks	Dark pool transaction datasets	Developed an enhanced transformer model for opaque market conditions	Dataset quality limitations; repeated listing suggests overlap
21	Ashtiani (2021)	Fraud detection in financial statements	Machine Learning and Data Mining	Financial statement data	Provided a systematic literature review for intelligent fraud detection	Focused on offline financial reports, not HFT or real-time data
22	Wang, Jiayi (2024)	Personalized search enhancement using AI	Hybrid Deep Learning and Cloud Computing	User interaction datasets	Proposed hybrid model improving search personalization	Domain not related to financial trading or anomaly detection
23	Wang, Shikai (2024)	Accelerating deep learning training	Distributed High-Performance Computing methods	Deep learning training datasets	Improved computational efficiency for training deep models	Focused on system acceleration, not anomaly detection
24	Ibikunle (2024)	Potential of machine learning in HFT analysis	Survey and conceptual insights	Various public and private datasets reviewed	Highlighted the opportunities ML offers in understanding HFT behaviors	Lacks experimental model validation and

						implementation examples
25	Zhang, Lu (2025)	Major issues in high-frequency financial data analysis	Survey and solution categorization	Multiple HFT datasets	Presented key challenges and survey of solutions in financial HFT data analysis	No experimental model or comparative benchmarks provided
26	Bello, Halima (2024)	Adaptive ML models for real-time fraud detection	Conceptual discussion of adaptive learning	Financial fraud environments (conceptual)	Addressed need for dynamic models in fraud prevention systems	Lack of experimental setups or real-world deployments
27	Basit (2024)	Quantum Variational Autoencoders for HFT anomaly detection	Quantum Machine Learning (QVAE)	Simulated high-frequency trading datasets	Proposed using quantum variational autoencoders to enhance anomaly detection	Early-stage; quantum computing feasibility remains uncertain
28	Nguyen (2021)	Forecasting and anomaly detection in supply chains	LSTM and LSTM Autoencoders	Supply chain time-series datasets	Demonstrated LSTM-based models improving anomaly detection and forecasting	Non-financial datasets; supply chain context, not HFT
29	Zhao, Zihan (2020)	Limit order book modeling for price prediction	Deep Convolutional Neural Networks (DeepLOB)	Public limit order book datasets	Achieved strong predictive performance using CNNs for order book data	Focused on price prediction, not explicitly on anomaly detection
30	Sirignano (2019)	Universal features of price formation using deep learning	Deep Neural Networks	Extensive limit order book data	Identified deep universal patterns in financial market behavior	Anomaly detection not the primary focus; descriptive analysis
31	Zhang, Ziyu (2020)	Spatio-temporal modeling of financial data	Temporal Graph Attention Networks (TGAT)	Financial transaction datasets	Improved modeling of complex spatio-temporal relationships	Scalability to extremely high-frequency real-time data untested

32	Zhu, Zhiwei (2020)	Anomaly detection in HFT using LSTM networks	LSTM-based anomaly detection	High-frequency trading datasets	Demonstrated LSTM networks effectively detecting anomalies in HFT	Real-time latency considerations not deeply addressed
----	--------------------	--	------------------------------	---------------------------------	---	---

III. Taxonomy of Anomaly Detection Methods

Anomaly detection methods in high-frequency financial data can be systematically classified into nine broad categories: statistical, proximity-based, clustering-based, classification-based, reconstruction-based, ensemble, deep learning-based, graph-based, and hybrid methods. Statistical methods assume that most data obey a known distribution and flag observations in low-probability regions—using methods such as Z-score analysis or histogram-based detection. The distance-based approaches, including k-nearest neighbors and Local Outlier Factor, identify anomalies as points located at fairly large distances or regions of significantly lower density than their neighbors. Clustering-based approaches—such as k-means and DBSCAN—identify outliers as points not belonging to any dense cluster but small sparse clusters. Classification-based techniques approach anomaly detection as a semi-supervised or supervised learning issue, where one trains models such as Support Vector Machines and One-Class SVMs to distinguish anomalous and normal instances.

Reconstruction-based methods learn dense representations of normal data through autoencoders or Principal Component Analysis and label points with extreme reconstruction error as anomalies. Ensemble-based methods provide robustness through the use of multiple detectors—Isolation Forest and feature-bagging are prominent examples—to identify diverse patterns of anomalies. Deep learning-based methods leverage deep neural networks (CNNs, RNNs such as LSTMs/GRUs, and Variational Autoencoders) to construct complex, non-linear relationships within high-dimensional unstructured data. Graph-based methods represent trading or transaction data as graphs and apply Graph Neural Networks or subgraph analysis to learn structural anomalies, which scalar features cannot. Finally, hybrid methods integrate two or more of the above paradigms—i.e., combining clustering with reconstruction—to benefit from each other's complementary strengths and assist in overcoming individual weaknesses. Table 2 shows summary of anomaly detection methods.

Table 2. Summary of Anomaly Detection Methods.

Category	Key Techniques	Data Type	Supervision	Advantages	Limitations
Statistical Methods	Z-score, Histogram-based	Numeric	Unsupervised	Simple, interpretable	Assumes specific data distribution
Proximity-Based Methods	k-Nearest Neighbors (k-NN), LOF	Numeric	Unsupervised	Flexible, non-parametric	Computationally expensive in large/high-D
Clustering-Based Methods	k-Means, DBSCAN	Numeric	Unsupervised	Detects arbitrarily shaped clusters	Sensitive to parameter choice
Classification-Based	SVM, One-Class SVM	Numeric	Supervised / Semi-Supervised	High accuracy with labels	Requires labeled data
Reconstruction-Based	Autoencoders, PCA	High-dimensional	Unsupervised	Captures complex patterns	May reconstruct anomalies as normal

Category	Key Techniques	Data Type	Supervision	Advantages	Limitations
Ensemble Methods	Isolation Forest, Feature Bagging	Various	Unsupervised	Robust, captures diverse patterns	Increased model complexity
Deep Learning-Based	CNNs, RNNs (LSTM/GRU), VAEs	Unstructured	Unsupervised	Models non-linear, hierarchical features	Data- and compute-intensive
Graph-Based Methods	Graph Neural Networks, Subgraph Analysis	Rational	Unsupervised / Semi-Supervised	Captures structural relationships	Graph construction and scaling challenges
Hybrid Methods	Combinations of above (e.g., clustering + autoencoder)	Mixed	Varies	Leverages complementary strengths	Design and tuning complexity

IV. Key Findings, Research Gaps and Proposed Directions

A. Higher-Level and Niche Domain Selection

The overarching research domain focuses on time-series analysis for financial systems, particularly in areas such as risk management, volatility analysis, and explainable artificial intelligence (XAI) for financial markets. At a more specialized level, niche domains have emerged, notably anomaly detection in high-frequency financial data for options markets, interpretable models for risk prediction, and the application of explainable AI techniques in volatility forecasting.

The intraday volatility prediction in options markets presents a highly promising research area, with the potential to integrate news feeds and sentiment analysis to enhance forecasting accuracy. This direction remains underexplored and offers substantial scope for future work.

B. Identified Research Gaps and Opportunities for Improvement

- Several critical gaps and improvement areas were identified across existing research studies:
- **Ensemble and Hybrid Modeling:** Very few studies have explored the combination of multiple machine learning paradigms (e.g., clustering, density estimation, reconstruction, forecasting) into ensemble or hybrid architectures. This remains a promising avenue to enhance model robustness.
 - **Limited Focus on Options Data:** The majority of studies concentrate on equities or futures; options markets, despite their growing importance, remain largely neglected in anomaly detection frameworks.
 - **False Positive and False Negative Challenges:** Many existing systems suffer from high false positive or false negative rates, necessitating the development of more precise detection models.
 - **Underutilization of Transformer-Based Attention:** Transformer architectures, particularly attention mechanisms, have not been sufficiently exploited for sequential financial data, especially in the context of anomaly detection.
 - **Insufficient Semi-Supervised Learning:** Semi-supervised approaches, which are critical when labeled anomaly data is scarce, are not extensively applied in the current literature.
 - **Minimal Use of Clustering Techniques:** Clustering methods, particularly for unsupervised anomaly detection, have been underutilized.
 - **Explainable AI (XAI) for Trustworthiness:** While anomaly detection models are advancing, their interpretability remains poor. Techniques such as SHAP and LIME could be deployed to explain model predictions, thereby improving trust and regulatory acceptance.

C. Metrics and Evaluation Criteria

Robust evaluation metrics are crucial for anomaly detection systems in high-frequency financial environments. The following metrics are recommended:

- **Classification Metrics:** ROC-AUC, Precision, Recall, and F1-Score to measure detection performance.
- **Imbalanced Dataset Metrics:** Matthews Correlation Coefficient (MCC) for assessing performance under data imbalance.
- **Operational Metrics:** Confusion matrix analysis and latency measurements for real-time deployment assessment.
- **Profitability Metrics:** Sharpe ratio calculation to evaluate profitability impact.
- **Portfolio Stability:** Volatility reduction measures in the managed portfolio.
- **Robustness:** Stress testing under extreme market conditions to assess model resilience.

D. Specific Findings and Observations

Through a synthesis of previous research and analysis, the following observations have been made:

- **Deep Learning Strengths:** Deep neural networks effectively learn complex, hierarchical data structures, crucial for identifying subtle market anomalies.
- **Autoencoders and Anomaly Detection:** Autoencoders, particularly in unsupervised or semi-supervised settings, are highly effective. LSTM Autoencoders outperform standard LSTM models in certain scenarios.
- **Simulated Fraud Data:** Due to the rarity of real-world manipulation cases, synthetic fraud datasets are often generated by injecting simulated anomalies into clean datasets.
- **Feature Engineering Importance:** Features such as price spreads, volatility measures, trading volume anomalies, and divergence between option and underlying asset prices are essential signals.
- **Options Market Specificity:** Metrics like implied volatility, option Greeks (delta, gamma, theta, vega), and order flow dynamics are critical in detecting anomalies unique to options markets.
- **Volatility Prediction Techniques:** GARCH models, SVMs, random forests, and LSTMs have demonstrated varying levels of success in volatility prediction tasks.
- **Explainability Need:** It is advisable to build inherently interpretable models rather than relying exclusively on post-hoc explainability techniques.
- **Handling Dimensionality:** Techniques such as feature selection (MRMR, CMIM), dimensionality reduction (PCA, t-SNE), and deep autoencoders are necessary to combat the curse of dimensionality.

E. Proposed Research Directions

Based on the identified gaps and findings, the following future directions are proposed:

- **Hybrid Architecture:** Develop models combining LSTM for sequence modeling, Transformers for attention mechanisms, and GANs for synthetic anomaly generation.
- **Multi-Paradigm Anomaly Detection:** Create ensemble frameworks that integrate clustering, density estimation, reconstruction, and forecasting models.
- **Options Market Focus:** Build dedicated anomaly detection systems for options data, considering OI spikes, implied volatility shifts, and unusual order flows.
- **Real-Time Processing Pipelines:** Leverage Apache Kafka for ingestion and Apache Flink or Spark Streaming for real-time anomaly detection.

- **Model Update Mechanisms:** Implement reinforcement learning or incremental learning strategies to enable continuous model adaptation without full retraining.
- **Explainable AI Integration:** Integrate SHAP and LIME explainability from model inception to enhance trustworthiness and facilitate regulatory compliance.
- **Stock Grouping Strategies:** Cluster stocks by industry (e.g., healthcare, FMCG) to capture sector-specific anomaly patterns.
- **Stress Testing Models:** Regularly subject detection models to extreme market simulation scenarios to ensure robustness.

Finally, it is important to emphasize that models should be rigorously validated across multiple datasets, as algorithm performance can vary significantly between different financial instruments and trading environments. No universal model fits all scenarios, highlighting the necessity of empirical validation tailored to specific application domains.

V. Research Motivation, Objectives, and Proposed Framework

A. Research Motivation

High-frequency financial data, particularly in options markets, frequently exhibit anomalous behaviors due to market microstructure noise, manipulative activities, and sudden price movements. Traditional anomaly detection methods, including statistical techniques such as Z-score analysis and machine learning models like isolation forests and one-class SVMs, struggle to handle the complex, high-dimensional, and nonstationary nature of such data. Moreover, deep learning approaches, although powerful, often function as opaque black boxes, limiting their applicability in regulated financial environments that require transparent decision-making.

Consequently, there exists a critical need to develop robust, adaptive, and explainable anomaly detection frameworks specifically tailored for high-frequency options data. Similarly, volatility forecasting models traditionally utilized in finance, such as GARCH and stochastic volatility models, often fail to capture nonlinear market dependencies or adapt effectively to rapidly evolving conditions. This research is motivated by the goal of advancing fraud detection capabilities, improving trading strategy optimization, and enhancing financial system resilience through interpretable, deep learning-based techniques.

B. Research Objectives

This study seeks to bridge existing gaps by pursuing multiple interconnected objectives. The first objective is to develop a deep learning-based anomaly detection framework leveraging autoencoders, attention mechanisms, and reinforcement learning strategies to identify irregularities within high-frequency options data streams. Secondly, the research emphasizes enhancing model interpretability by incorporating explainable AI (XAI) techniques, particularly SHAP (SHapley Additive Explanations), LIME, and attention visualization, ensuring that model outputs can be easily understood and validated.

Another key objective involves the integration of domain-specific feature engineering, employing financial attributes such as implied volatility surfaces, bid-ask spread patterns, and the Greeks (delta, gamma, theta, vega) to capture nuanced anomalies that generic features might overlook. The study also aims to address scalability challenges by employing transformer-based architectures optimized for low-latency inference. Finally, a structured benchmarking and evaluation framework will be established to validate model performance across diverse market regimes, emphasizing robustness, adaptability, and explainability.

The key features extracted for anomaly detection in high-frequency options data, along with their financial relevance, are summarized in Table 2.

C. Proposed Methodological Framework

The proposed research framework encompasses a structured multi-phase methodology. Data will be collected from high-frequency options trading sources, including NSE, CBOE, and Upstox APIs, capturing rich features such as implied and historical volatility, options Greeks, bid-ask spread dynamics, and order flow imbalance. The datasets will be preprocessed to handle missing values, normalize key attributes, and segment trading sessions based on market regimes.

Model development will initially involve establishing baseline architectures, including standard autoencoders, variational autoencoders (VAEs), and GAN-based anomaly detectors. These models will be used as comparative benchmarks. Sophisticated model building will incorporate transformer-based architectures with the ability to model sequential dependencies and attention-based prioritization of features. Reinforcement learning algorithms will be embedded in order to permit dynamic modification of anomaly detection thresholds according to changing market conditions. In addition, regime-switching neural networks will be investigated in order to modulate model behavior according to real-time changes in volatility or liquidity.

Explainability will be infused at inception, where SHAP and LIME frameworks are used to spot significant features determining model outputs. Attention mechanisms from transformer-based models will be mapped to help enhance interpretability as well as extract actionable insights for market anomalies.

Performance will be measured through a blend of standard classification metrics like Precision, Recall, F1-Score, and ROC-AUC, as well as financial performance measures like Sharpe Ratio gains and Matthews Correlation Coefficient (MCC) in imbalanced cases. Comparative testing against standard anomaly detection methods like Z-score analysis, Isolation Forests, and one-class SVMs will be performed. Stress testing under volatile and stressed market conditions will be an integral part of the benchmarking process to ensure model robustness.

As far as deployment of real-time systems is concerned, the research suggests using Apache Kafka for ingesting streaming data and Apache Flink or Spark Streaming for scalable low-latency processing pipelines.

Continuous learning strategies, incorporating mini-batch incremental updates and reinforcement learning enhancements, will ensure that models adapt dynamically without the need for full retraining. Data augmentation using GANs will address the imbalance in fraud versus legitimate data cases by synthesizing realistic anomalous patterns.

Figure 1 describes the proposed architecture diagram of the system.

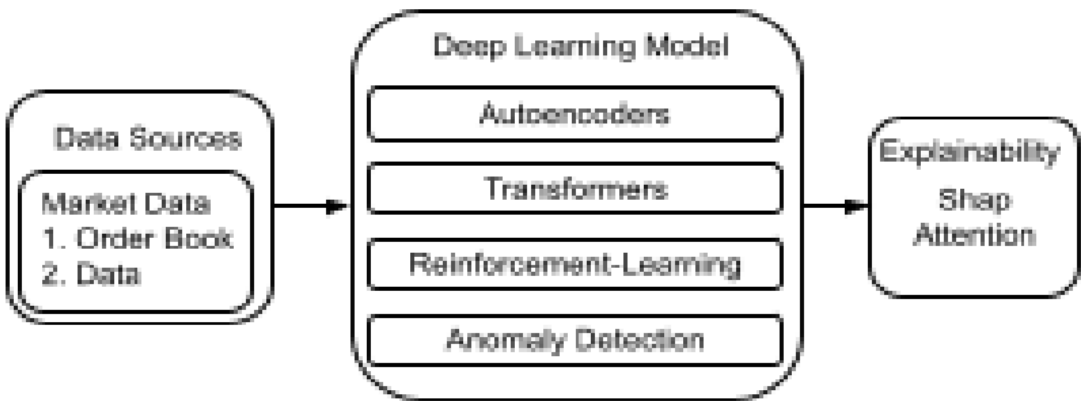


Figure 1. Proposed Architecture Diagram.

The evaluation metrics selected for assessing the anomaly detection framework, including both classification and financial impact measures, are outlined in Table 3.

Table 3. Feature Set for Anomaly Detection in Options Data.

Feature	Description	Relevance to Anomaly Detection
Implied Volatility	Market's expectation of future volatility derived from option prices.	Identifies sudden changes in perceived market risk.
Historical Volatility	Past realized volatility of the underlying asset.	Detects discrepancies between past and expected volatility.
Delta	Sensitivity of option price to changes in the underlying asset price.	Captures unusual hedging activity or manipulation.
Gamma	Rate of change of Delta relative to the underlying asset price.	Highlights nonlinear price dynamics during anomalies.
Theta	Time decay of an option's value.	Observes pricing anomalies related to time-value erosion.
Vega	Sensitivity of option price to changes in volatility.	Identifies volatility-driven anomalies.
Bid-Ask Spread	Difference between highest buying and lowest selling prices.	Captures liquidity disruptions or artificial widening.
Order Flow Imbalance	Net difference between buy and sell orders.	Detects spoofing, momentum ignition, and liquidity shifts.
Best Bid and Ask Prices	Highest bid and lowest ask prices available.	Useful for detecting order book spoofing or layering.
Open Interests	Number of open contracts outstanding.	Identifies abnormal buildup or unwinding positions.

Table 4. Evaluation Metrics and Their Purpose.

Metric	Purpose	Notes on Importance
ROC-AUC	Measures model's ability to distinguish between classes.	Robust to class imbalance, good for anomaly detection tasks.

Precision	Proportion of true positives among predicted positives.	Important when false positives are costly.
Recall	Proportion of true positives detected among all actual positives.	Important to catch rare but critical anomalies.
F1-Score	Harmonic mean of Precision and Recall.	Balances false positives and false negatives.
Matthews Correlation Coefficient (MCC)	Measures quality of binary classifications, even with imbalanced classes.	Critical for rare-event anomaly detection.
Sharpe Ratio Impact	Evaluates improvement in portfolio profitability after anomaly detection.	Links model accuracy with trading strategy profitability.
Latency Measurement	Time taken to detect anomalies after occurrence.	Essential for real-time financial market applications.
Stress Testing Under Market Regimes	Testing model under extreme volatility and low liquidity conditions.	Validates robustness and practical usability of the model.

D. Expected Contributions

Through this research, several key contributions are anticipated. Firstly, a novel anomaly detection framework tailored specifically for high-frequency options data will be developed, advancing the field beyond current equity-focused models. Secondly, the integration of XAI techniques within deep learning architectures will address the critical challenge of model transparency in financial applications. Thirdly, a benchmarking methodology for evaluating both detection accuracy and interpretability of deep financial models will be proposed. Lastly, practical insights into market microstructure behaviors—particularly regarding volatility and liquidity disruptions in options markets—will be generated, informing both academic research and industry practices.

VI. Conclusion

Anomaly detection in high-frequency financial data has emerged as a crucial area of research, especially in the context of options markets where market microstructure complexities and non-linear dynamics are prominent. This review paper systematically explored the advancements in anomaly detection techniques, emphasizing the application of deep learning models tailored to the unique challenges posed by high-frequency trading environments. Through an in-depth analysis of existing

literature, it was observed that while significant progress has been made using deep architectures such as LSTM networks, graph neural networks, and transformer models, several critical gaps remain.

Specifically, the lack of real-time adaptability, limited interpretability, and insufficient domain-specific feature engineering were recurrent limitations across multiple studies. Moreover, relatively few contributions have focused on the options segment, despite its growing importance in modern financial systems. The necessity for hybrid models combining statistical methods and deep learning, the integration of explainable AI techniques, and the adoption of streaming data architectures were identified as key directions for future research.

To address these gaps, this review proposes a comprehensive framework leveraging advanced deep learning models, reinforcement learning for dynamic adaptation, and explainable AI methodologies for enhanced transparency. The incorporation of real-time processing pipelines, including tools such as Apache Kafka and Apache Flink, was also highlighted to meet the operational demands of high-frequency financial markets.

In conclusion, the intersection of deep learning, real-time anomaly detection, and explainable AI in high-frequency financial data represents a promising yet underexplored frontier. Advancing research in this area holds the potential to significantly enhance market integrity, support regulatory oversight, and enable the development of robust, trustworthy algorithmic trading systems. Future work should emphasize not only achieving high detection accuracy but also ensuring system interpretability, scalability, and resilience across diverse market regimes.

VII. Glossary of Terms

Anomaly Detection: The act of detecting data points, events, or observations that are very different from the norm or expected behavior in a dataset. In time-series situations, anomalies can signal faults, fraud, or new patterns to be investigated.

High-Frequency Trading (HFT): A form of algorithmic trading characterized by high speeds, turnover, and order-to-trade ratios, employing sophisticated algorithms and co-location to execute high volumes of orders in fractions of a second.

Long Short-Term Memory (LSTM): A sophisticated recurrent neural network structure that was suggested to minimize vanishing-gradient issues using memory cells and gates, so long-range temporal dependencies in sequential data can be modeled.

Transformer (Deep Learning Architecture): A deep learning model based only on multi-head self-attention mechanisms, without recurrence for parallel sequence data processing and improved global relationships capture.

Graph Neural Network (GNN): A family of neural networks that operate on graph structures, learning node, edge, or graph representations by aggregating and processing information along graph connectivity.

Autoencoder: An unsupervised neural network that maps input data onto a compressed (encoded) low-dimensional representation and back to the original space (decoding), with reconstruction error being used for purposes like anomaly detection.

Reinforcement Learning (RL): A framework in which an agent acts in a world, receiving rewards or penalties, and learns to act so as to maximize total reward over time by trial-and-error exploration.

SHAP (SHapley Additive exPlanations): An interpretability technique based on game-theoretic Shapley values that assigns an importance score to every feature for a particular prediction, enabling transparent, locally coherent explanations.

LIME (Local Interpretable Model-Agnostic Explanations): A model-agnostic method that locally approximates any black-box predictor with an explainable model (e.g., a simple decision tree), providing human-interpretable explanations for individual predictions C3 AI.

Limit Order Book (LOB): A current record of all outstanding limit buy and sell orders for a security, ordered by price tier, which reflects market depth and is at the center of many high-frequency trading techniques Wikipedia.

Attention Mechanism: An algorithm that estimates a weighted average of input parts, allowing models to focus on the most salient features of a sequence for prediction or reconstruction.

One-Class Classification: A learning situation in which a model is trained exclusively on examples of the "normal" class and must decide whether new cases fall outside this standard, such as One-Class SVMs.

Local Outlier Factor (LOF): A density-based anomaly detection method that calculates how far away a point is from the locally determined k-nearest neighbors.

Empirical Mode Decomposition (EMD): A signal processing technique that separates a non-stationary time series into intrinsic mode functions from local minima and maxima, useful for isolating high-frequency components before anomaly detection.

Kernel Density Estimation (KDE): A method of approximating the probability density function of a random variable by smoothing kernel functions (e.g., Gaussian) around every data point.

Federated Learning: A privacy-preserving distributed learning framework where multiple clients (e.g., institutions) collaboratively train a shared model without exposing their data from devices or servers.

References

1. Brugman, S. R. D. "The development of a real-time monitoring system for fatigue detection on truckers." Bachelor's thesis, University of Twente, 2022.
2. Veryzhenko, Iryna, Nohade Nasrallah, and Henri Garcia. "Detecting spoofing in high frequency trading using machine learning techniques."
3. Rizvi, Baqar, Ammar Belatreche, and Ahmed Bouridane. "Stock Price Manipulation Detection using Empirical Mode Decomposition based Kernel Density Estimation Clustering Method." (2018).
4. Zhao, Hang, Yujing Wang, Juanyong Duan, Congrui Huang, Defu Cao, Yunhai Tong, Bixiong Xu, Jing Bai, Jie Tong, and Qi Zhang. "Multivariate time-series anomaly detection via graph attention network." In 2020 IEEE international conference on data mining (ICDM), pp. 841-850. IEEE, 2020.
5. Mejri, Nesryne, Laura Lopez-Fuentes, Kankana Roy, Pavel Chernakov, Enjie Ghorbel, and Djamila Aouada. "Unsupervised anomaly detection in time-series: An extensive evaluation and analysis of state-of-the-art methods." Expert Systems with Applications (2024): 124922.
6. Chen, Yuexing, Maoxi Li, Mengying Shu, Wenyu Bi, and Siwei Xia. "Multi-modal Market Manipulation Detection in High-Frequency Trading Using Graph Neural Networks." Journal of Industrial Engineering and Applied Science 2, no. 6 (2024): 111-120.
7. Xi, Yue, Yining Zhang, and Hanqing Zhang. "Real-time Multimodal Route Optimization and Anomaly Detection for Cross-border Logistics Using Deep Reinforcement Learning." Academia Nexus Journal 3, no. 3 (2024).
8. Guanghe, Cao, Shuaiqi Zheng, Yibang Liu, and Maoxi Li. "Real-time anomaly detection in dark pool trading using enhanced transformer networks." Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online) 3, no. 4 (2024): 320-329.
9. Cao, Guanghe, Yitian Zhang, Qi Lou, and Gaike Wang. "Optimization of High-Frequency Trading Strategies Using Deep Reinforcement Learning." Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023 6, no. 1 (2024): 230-257.
10. Li, Maoxi, Mengying Shu, and Tianyu Lu. "Anomaly Pattern Detection in High-Frequency Trading Using Graph Neural Networks." Journal of Industrial Engineering and Applied Science 2, no. 6 (2024): 77-85.
11. Alaminos, David, M. Belén Salas, and Antonio Partal-Ureña. "Hybrid ARMA-GARCH-Neural Networks for intraday strategy exploration in high-frequency trading." Pattern Recognition 148 (2024): 110139.
12. Shanmuganathan, V., and Annamalai Suresh. "Markov enhanced I-LSTM approach for effective anomaly detection for time series sensor data." International Journal of Intelligent Networks 5 (2024): 154-160.

13. Bello, Halima Oluwabunmi, Adebimpe Bolatito Ige, and Maxwell Nana Ameyaw. "Deep learning in high-frequency trading: conceptual challenges and solutions for real-time fraud detection." *World Journal of Advanced Engineering Technology and Sciences* 12, no. 02 (2024): 035-046.
14. Li, Lin, Yitian Zhang, Jiayi Wang, and Ke Xiong. "Deep Learning-Based Network Traffic Anomaly Detection: A Study in IoT Environments." (2024).
15. M. Jin et al., "A Survey on Graph Neural Networks for Time Series: Forecasting, Classification, Imputation, and Anomaly Detection," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, pp. 10466-10485, Dec. 2024, doi: 10.1109/TPAMI.2024.3443141.
16. Haq, Ijaz Ul, Byung Suk Lee, Donna M. Rizzo, and Julia N. Perdrial. "An automated machine learning approach for detecting anomalous peak patterns in time series data from a research watershed in the northeastern United States critical zone." *Machine Learning with Applications* 16 (2024): 100543.
17. Bakumenko, Alexander, and Ahmed Elragal. "Detecting anomalies in financial data using machine learning algorithms." *Systems* 10, no. 5 (2022): 130.
18. Poutré, Cédric, Didier Chételat, and Manuel Morales. "Deep unsupervised anomaly detection in high-frequency markets." *The Journal of Finance and Data Science* 10 (2024): 100129.
19. Yang, Mingxuan, Decheng Huang, Weixiang Wan, and Meizhizi Jin. "Federated learning for privacy-preserving medical data sharing in drug development." *Applied and Computational Engineering* 108 (2024): 7-13.
20. Guanghe, Cao, Shuaiqi Zheng, Yibang Liu, and Maoxi Li. "Real-time anomaly detection in dark pool trading using enhanced transformer networks." *Journal of Knowledge Learning and Science Technology* 3, no. 4 (2024): 320-329.
21. Ashtiani, Matin N., and Bijan Raahemi. "Intelligent fraud detection in financial statements using machine learning and data mining: a systematic literature review." *Ieee Access* 10 (2021): 72504-72525.
22. Wang, Jiayi, Tianyu Lu, Lin Li, and Decheng Huang. "Enhancing personalized search with ai: a hybrid approach integrating deep learning and cloud computing." *Journal of Advanced Computing Systems* 4, no. 10 (2024): 1-13.
23. Wang, Shikai, Haotian Zheng, Xin Wen, and Shang Fu. "Distributed high-performance computing methods for accelerating deep learning training." *Journal of Knowledge Learning and Science Technology* ISSN: 2959-6386 (online) 3, no. 3 (2024): 108-126.
24. Ibikunle, Gbenga, Ben Moews, Dmitriy Muravyev, and Khaladdin Rzayev. "Can machine learning unlock new insights into high-frequency trading?." *arXiv preprint arXiv:2405.08101* (2024).
25. Zhang, Lu, and Lei Hua. "Major Issues in High-Frequency Financial Data Analysis: A Survey of Solutions." *Mathematics* 13, no. 3 (2025): 347.
26. Bello, Halima Oluwabunmi, Adebimpe Bolatito Ige, and Maxwell Nana Ameyaw. "Adaptive machine learning models: concepts for real-time financial fraud prevention in dynamic environments." *World Journal of Advanced Engineering Technology and Sciences* 12, no. 02 (2024): 021-034.
27. Basit, Jamshaid, Danish Hanif, and Madiha Arshad. "Quantum Variational Autoencoders for Predictive Analytics in High Frequency Trading Enhancing Market Anomaly Detection." *International Journal of Emerging Multidisciplinaries: Computer Science & Artificial Intelligence* 3, no. 1.
28. Nguyen, Huu Du, Kim Phuc Tran, Sébastien Thomassey, and Moez 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.
29. Zhao, Zihan, Longfei Li, Qingyun Wu, and Liuyi Yao. "DeepLOB: Deep Convolutional Neural Networks for Limit Order Books." *IEEE Transactions on Signal Processing* 68 (2020): 1441-1452.
30. Sirignano, Justin, and Rama Cont. "Universal Features of Price Formation in Financial Markets: Perspectives from Deep Learning." *Quantitative Finance* 19, no. 9 (2019): 1449-1459.
31. Zhang, Ziyu, Yichuan Charlie Hu, and Srinivasan Parthasarathy. "Spatio-Temporal Modeling of Financial Data via Temporal Graph Attention Networks." *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM)*, 2020.
32. Zhu, Zhiwei Steven, and Timothy Chan. "Anomaly Detection in High-Frequency Trading Data Using LSTM Networks." *Journal of Financial Data Science* 2, no. 1 (2020): 55-69.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.