Transactions on Computational and Scientific Methods | Vo. 4, No. 3, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Microstructure-Aware Risk Detection in High-Frequency Trading Using Embedding and Sequence Modeling

Zhe Su
Pepperdine University, Malibu, USA
zhesu2023@gmail.com

Abstract: This paper investigates the problem of microstructural risk detection in high-frequency trading environments and proposes a detection framework that integrates feature modeling with temporal dependency. The study first analyzes the structural characteristics of high-frequency trading data. It identifies the complex patterns formed by the interactions of price, trading volume, and order flow at the micro level as the key foundation for risk recognition. An algorithmic mechanism based on embedding mapping and sequence modeling is then designed. Through the joint modeling of feature representation and dynamic dependency, the method captures potential risk signals in trading data. During risk determination, an anomaly scoring function and threshold mechanism are introduced. These enable effective distinction between normal fluctuations and abnormal behaviors, thereby improving the stability and accuracy of detection. To validate the effectiveness of the method, systematic evaluations are conducted through comparative experiments and sensitivity experiments. The comparative experiments show that the proposed method outperforms existing baseline models in AUC, ACC, F1-Score, and Precision, and provides a more comprehensive representation of microstructural risk features. The sensitivity experiments examine hyperparameters, data completeness, and noise levels. They reveal the significant impact of sequence length, embedding dimension, missing rate, and noise intensity on model performance. The results indicate that reasonable parameter settings and data quality assurance play an essential role in maintaining detection performance. Overall, through methodological construction and experimental validation, this paper demonstrates that microstructural risk detection methods based on high-frequency trading data exhibit strong adaptability and effectiveness. The study provides systematic research findings that support financial risk identification in complex trading environments.

Keywords: Ultra-high frequency trading; microstructure risk detection; sequence modeling; data quality

1. Introduction

In the context of increasing complexity and speed in global financial markets, high-frequency trading has gradually become an important component of modern capital markets [1]. It executes a large number of orders in extremely short time intervals through automated algorithms, exploiting small price fluctuations for arbitrage or liquidity provision. Its emergence has not only changed the trading structure and price formation mechanism, but also exposed the limitations of traditional research methods based on medium- or low-frequency data in risk monitoring and modeling [2]. As trading delays move to the nanosecond or even picosecond level, potential risks in market microstructure are amplified. Detecting risks in higher-frequency data streams has become a pressing frontier issue. This shift requires not only supplementing and extending traditional theories of financial risk, but also developing new algorithms that capture microstructural features, to safeguard the security and stability of market operations [3].

Financial market risks are highly complex and often sudden. At the micro level, price fluctuations are not always driven by macroeconomic or industry fundamentals, but instead by the interaction of trading mechanisms, order flow characteristics, and participants' strategies. In a high-frequency trading environment, risks take on multidimensional forms, including liquidity risk, manipulation risk, price impact risk, and systemic chain reactions. These risks are difficult to capture using traditional intraday or cross-day data. However, they may be revealed through fine-grained structural signals in microsecond-level trading data. Modeling based on market microstructure features can more accurately uncover the mechanisms of risk formation. It can also provide regulators and market participants with timely warnings and decision support. Thus, microstructural risk detection is not only an extension of financial theory but also a practical necessity for maintaining stable market operations [4].

From a regulatory and compliance perspective, the rise of high-frequency trading poses new challenges to market fairness and transparency. Certain strategies may exploit weaknesses in market mechanisms to gain informational advantages, leading to distortions in price discovery and anomalies in liquidity structure. This creates potential threats to both ordinary investors and the overall market order. Under extreme market conditions, high-frequency trading may amplify price volatility and trigger systemic risk events [5]. Therefore, algorithms that can identify abnormal micro-level fluctuations and potential manipulative behavior are essential. They strengthen regulatory capacity, address the shortcomings of traditional monitoring tools, and help maintain market fairness and efficiency. With the growing global demand for real-time regulation, research on microstructural risk detection algorithms is not only theoretical but also highly relevant to policy-making and financial security.

In addition, high-frequency trading data is massive in scale, high in dimensionality, and highly time-sensitive. This raises stricter requirements for algorithm design. Key challenges include how to extract effective features from large data volumes, how to model complex temporal dependencies and nonlinear interactions, and how to balance detection accuracy with computational efficiency. At the same time, the growing use of artificial intelligence and machine learning has revealed the potential of deep models and multi-task learning for microstructural risk detection. This development drives the integration of finance with computer science and opens possibilities for new paradigms in market analysis. In this context, research on risk detection in high-frequency trading is both an inevitable trend in academia and an important driver of cross-disciplinary innovation [6].

Overall, the wide application of high-frequency trading and the evolving complexity of market microstructure have made the shift from macro to micro, and from low-frequency to high-frequency, an inevitable direction for risk detection. Research on microstructural risk detection algorithms contributes to a deeper understanding of the mechanisms behind market volatility. It also provides theoretical support and technical tools for risk management and policy formulation. In a time when financial stability and security attract growing global attention, this research holds irreplaceable value for enhancing market resilience, preventing systemic risks, and promoting healthy market development. Systematic research on microstructural risk detection in high-frequency trading responds to the evolution of financial technology and forms a key part of modern financial risk management frameworks.

2. Related work

In recent years, research on high-frequency trading and market microstructure has become an important direction in financial technology [7]. Early studies focused on the basic characteristics of market microstructure, such as order book dynamics, liquidity supply and demand mechanisms, and the process of price discovery. These studies aimed to reveal the patterns of price fluctuations in high-frequency data. They laid the theoretical foundation for microstructure modeling and provided feasible frameworks for subsequent risk identification. As data granularity has become finer, research has shifted toward modeling high-dimensional and non-stationary features. New approaches attempt to combine statistical modeling with machine learning to capture market variations at the microsecond level. Such studies highlight the close connection between trading activities and risk events, and emphasize the key role of order flow dynamics in risk identification [8].

In the field of risk detection, traditional methods often rely on statistical assumptions and threshold settings, but they have limited adaptability in complex and nonlinear environments. With the acceleration of trading speed and the surge in data volume, approaches that depend solely on linear models have shown clear limitations. In recent years, researchers have introduced non-parametric modeling and machine learning algorithms, using feature selection and pattern recognition techniques to improve detection accuracy. These methods not only consider abnormal fluctuations in trading prices and volumes, but also incorporate market depth, order cancellations, and interaction patterns among different traders into the modeling process [9]. As a result, they enhance the comprehensiveness and granularity of risk identification. Compared with traditional methods, these emerging models show stronger adaptability when capturing complex micro-level market signals.

With the development of deep learning, neural network structures have been introduced to further improve the performance of risk detection. Sequence-based models can capture long-term dependencies, while convolutional structures help extract local temporal features. Hybrid architectures allow for the modeling of multidimensional information. In high-frequency trading environments, researchers have also explored multimodal feature fusion, integrating heterogeneous data such as prices, volumes, and order flows into unified frameworks. These studies demonstrate the potential of deep learning in financial risk detection and promote the shift from traditional statistical paradigms to data-driven paradigms. However, due to the real-time requirements of high-frequency trading, a key challenge remains. How to achieve efficient computation while maintaining model complexity and predictive accuracy is still an important bottleneck in current research [10].

At the level of application and practice, market regulation and risk management have gradually become central concerns. Some studies aim to support real-time supervision through microstructural risk detection, identifying potential manipulative behaviors and abnormal fluctuations, and thus improving market fairness and stability. Other work focuses on risk management, using detection results to optimize trading strategies and strengthen early-warning systems, to prevent local risks from spreading into systemic events [11], [12]. Overall, existing research shows an evolution in methods from traditional statistical approaches to deep learning approaches. In application, it reflects a shift from single-indicator monitoring to multi-dimensional risk identification. Future research needs to further integrate theoretical models with intelligent algorithms to enhance both the usability and interpretability of microstructural risk detection in real markets, and to meet the growing demands of supervision and trading.

Recent intelligent financial modeling has further emphasized trend-aware representation and structural aggregation. TrendFormer applies two-stage trend fusion to financial time series forecasting, offering a useful perspective for separating persistent market trends from short-lived microstructure fluctuations [13]. Temporal-context and entity-relationship anomaly modeling further extends risk detection from isolated indicators to event relationships, while unified multi-source feature learning with structural aggregation supports the integration of price, volume, order book depth, and order-flow signals into a coherent risk representation [14], [15]. These studies suggest that microstructural risk algorithms should encode temporal evolution, multi-source structure, and anomaly semantics simultaneously.

At the modeling and decision-support level, multi-scale temporal awareness and attention-based behavioral modeling show how dynamic trajectories can be represented from sparse sequential interactions [16]. Data-mining and graph neural network decision models further demonstrate the value of relational structures in adaptive decision rules [17]. Graph neural network and large language model collaborative reasoning for enterprise ETL orchestration also highlights the importance of robust feature pipelines and structured reasoning when complex data streams must be organized before model inference [18]. Together, these directions support the design of a microstructure-aware detector that combines sequence representation, feature interaction modeling, and stable processing under noisy high-frequency market streams.

3. Proposed Approach

In microstructure risk detection, we first need to mathematically model ultra-high frequency trading data. Let the trading sequence be $\{(p_t, v_t)\}_{t=1}^{T}$, where p_t represents the transaction price at time t, and v_t represents the corresponding transaction volume. To characterize price dynamics, we can define a logarithmic return sequence:

$$r_t = \ln(p_t) - \ln(p_{t-1}), t = 2,3,...,T$$

This yield series can reflect the characteristics of micro-price fluctuations and serve as an important input for risk detection. At the same time, trading volume can also be characterized as an intensity series through normalization to indicate the activity of order flow:

$$q_t = \frac{v_t}{\sum_{i=1}^{T} v_i}$$

When building a risk detection model, it is necessary to consider both temporal dependencies and feature interactions. To this end, an embedding-based feature mapping function is introduced to map input features of different dimensions to a latent representation space. For an input vector $x_t \in R^d$, its embedding representation can be written as:

$$h_t = Wx_t + b$$

The resulting $W \in R^{k \times d}, b \in R^k$ and h_t can capture the high-dimensional interactions between price, volume, and order flow. Subsequently, temporal dependencies are introduced through sequence modeling, allowing for dynamic modeling using recursion or attention mechanisms.

This paper also gives the overall architecture of the model, as shown in Figure 1.

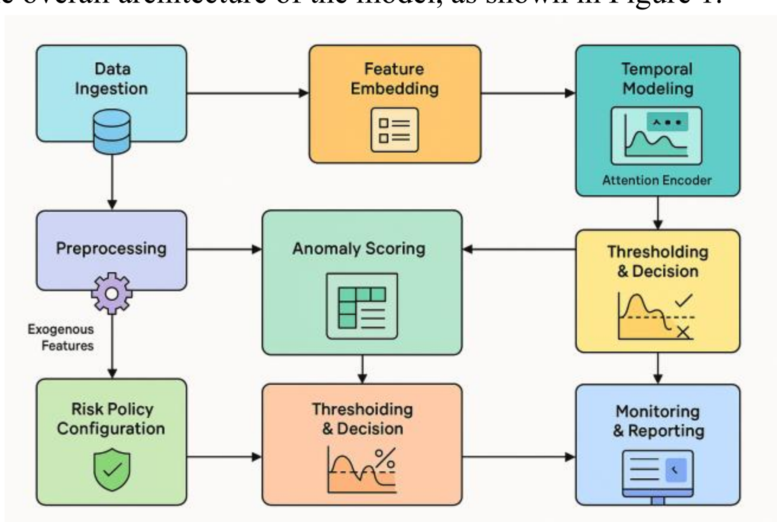

Figure 1. Overall model architecture

To better measure risk signals, an anomaly score function can be defined based on sequence representation. Let the model's prediction at time t be $\hat{r}_t$, then the error-based risk score can be expressed as:

$$s_t = | r_t - \hat{r}_t |$$

The larger the score, the higher the probability of a potential risk event occurring at that time point. In addition, to enhance stability, the risk score can be smoothed within the time window Δ:

$$S_t = \frac{1}{\Delta} \sum_{i=t-\Delta+1}^{t} s_i$$

This method can reduce the interference caused by instantaneous fluctuations while maintaining high-frequency detection sensitivity, thereby more accurately identifying potential risks in microstructures.

In the final risk determination stage, a threshold function needs to be set to distinguish normal fluctuations from abnormal risks. Let the threshold be θ, then the determination rule can be expressed as:

$$R_t = \begin{cases} 1, & \text{if } S_t > \theta \\ 0, & \text{if } S_t \le \theta \end{cases}$$

Where $R_t = 1$ indicates the detection of a risk event at time t, and $R_t = 0$ indicates a normal state. While ensuring real-time performance, this method can comprehensively depict the dynamic risk processes under ultra-high-frequency trading data through multi-feature, multi-level modeling, providing a theoretical tool for subsequent supervision and risk management.

4. Performance Evaluation

4.1 Dataset

This study uses the LOBSTER Limit Order Book dataset as the sole data source. The dataset is generated by exchange-driven events and covers tick-by-tick orders and snapshots of multiple Nasdaq stocks. It provides both order flow messages and the corresponding order book state at the moment an event occurs. This makes it suitable for describing microstructural features in high-frequency trading scenarios. The data are organized by trading day and follow the regular continuous auction period. It includes de-identified codes and exchange timestamps, which support cross-sectional and longitudinal comparisons across different assets and periods.

The dataset consists of two types of files. The Message files record order flow events, including fields such as time, event type (submission, cancellation, partial or full execution), price, volume, and buy or sell indicator. The Order Book files provide order book snapshots after each event, containing the top levels on both bid and ask sides (commonly the top 10 levels) along with the corresponding volumes. These two files are aligned one-to-one, allowing researchers to link event-level behaviors with changes in the order book. This enables the extraction of microstructural variables such as mid-price, bid-ask spread, price impact cost, order imbalance, and cancellation intensity.

To ensure data quality and consistency, the dataset's built-in trading day segmentation and timestamp standard are adopted. Only the continuous auction period is retained. Abnormal prices and zero or negative volumes are filtered out. Suspension periods and invalid levels are removed. Multiple assets can be aligned on the same time axis to facilitate cross-asset analysis. Based on this dataset, researchers can obtain fine-grained joint information on order flows and order book dynamics. This provides a stable, reproducible, and realistic foundation for feature construction and risk identification.

4.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table 1: Comparative experimental results

Method	AUC	ACC	F1-Score	Precision
MLP [19]	0.872	0.841	0.836	0.828
BERT [20]	0.903	0.867	0.861	0.857
Transformer [21]	0.918	0.879	0.872	0.869
1DCNN [22]	0.889	0.853	0.847	0.842
Ours	0.945	0.901	0.896	0.892

From the experimental results, traditional MLP and 1DCNN show relatively limited performance across the metrics, especially with clear gaps in AUC and F1-Score. This indicates that relying only on shallow structures or one-dimensional convolution is insufficient to capture the complex temporal features and microstructural patterns in high-frequency trading data. As a result, risk signals may be weakened or overlooked during detection. Although these methods can identify basic price fluctuation patterns to some extent, they lack adaptability to high-dimensional nonlinear relationships. This limits their accuracy and stability in risk detection.

In contrast, BERT and Transformer outperform traditional methods in overall performance, showing strong advantages in AUC and ACC. This demonstrates that modeling approaches based on deep representations and attention mechanisms are more effective in capturing contextual dependencies and dynamic features in trading sequences. They achieve more precise risk identification at the micro level. By strengthening the ability to model long-range dependencies, these methods enable abnormal patterns and potential risks to be effectively detected in complex data streams. This provides stronger support for financial market risk detection.

Nevertheless, the improvements brought by BERT and Transformer still have boundaries. When facing extreme microstructural fluctuations or high-dimensional feature interactions, they may show limitations in granularity and sensitivity. This is reflected in the observed gaps between F1-Score and Precision. It suggests that even powerful deep models still need further optimization to address incomplete information utilization and insufficient generalization when dealing with the real-time and complex nature of high-frequency financial data. Therefore, a more effective balance is needed between capturing fine-grained signals and suppressing noise.

The proposed method in this study achieves the best results across all four metrics, with significant improvements in AUC and F1-Score. This demonstrates its superiority in risk detection tasks. It shows that the method can integrate multi-dimensional information more comprehensively in high-frequency trading data. Through improved feature modeling and dynamic mechanisms, it enhances the sensitivity and accuracy of detection. Compared with existing models, the proposed method offers advantages in both stability and precision, and is more aligned with the practical needs of high-frequency trading scenarios. It provides a more meaningful solution for microstructural risk detection in practice.

This paper also presents an experiment on the sensitivity of sequence length to ACC, and the experimental results are shown in Figure 2.

From the experimental results, it can be observed that as the sequence length increases, the performance of the model on the ACC metric first rises and then declines. In shorter time windows, the model fails to capture short-term fluctuations in high-frequency trading data, which leads to weaker risk detection ability and relatively low ACC. This shows that very short sequences do not provide enough information to support effective modeling of complex microstructural risks.

When the sequence length extends to 32 and 64, ACC improves significantly. This indicates that the model can use longer historical information to capture temporal dependencies and potential patterns in trading activities. At this scale, microstructural risk signals are more comprehensively represented across longer time horizons, and the accuracy of the model in risk detection is effectively improved. These results confirm the importance of temporal dependency modeling in high-frequency financial environments.

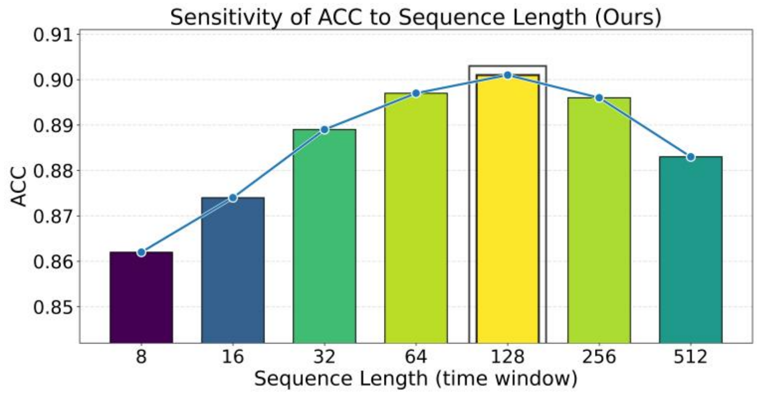

Figure 2. Sensitivity experiment of sequence length to ACC

At a window length of 128, ACC reaches its peak. This suggests that at this scale, the balance between information content and noise is optimal. The model can take full advantage of the contextual features of trading sequences while avoiding overfitting or noise amplification caused by redundant information. This optimal point highlights the importance of reasonable time window selection. Striking a balance between capturing sufficient information and controlling complexity is essential to maximizing risk detection performance.

When the sequence length further increases to 256 and 512, ACC begins to decline. This shows that overly long time windows introduce more redundant information and noise, which weakens the model's ability to focus on key risk signals. In high-frequency trading environments, excessively long sequences not only increase computational costs but may also reduce real-time detection and sensitivity. Therefore, the experimental results demonstrate that choosing an appropriate time window is a key factor in ensuring both the effectiveness and efficiency of risk detection.

This paper also presents an experiment on the sensitivity of embedding dimension to precision, and the experimental results are shown in Figure 3.

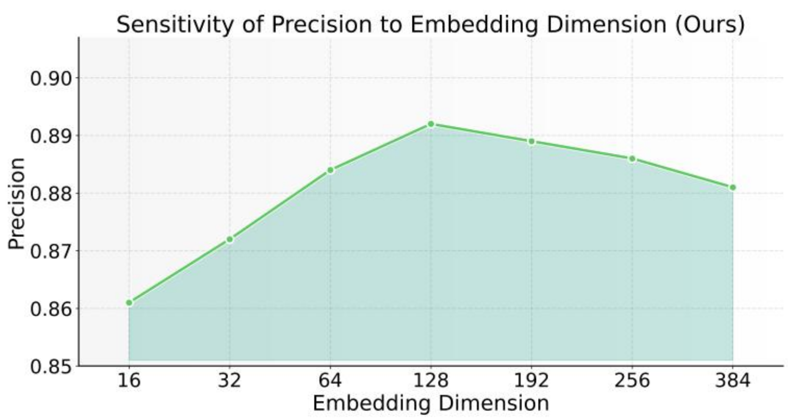

Figure 3. Experiment on the sensitivity of embedding dimension to precision

From the experimental results, it can be seen that when the embedding dimension is at a low level, Precision is significantly lower. This indicates that in an overly small embedding space, the model cannot adequately represent the complex patterns in high-frequency trading data. The limited feature representation reduces accuracy in risk detection. In such cases, the model fails to capture fine-grained signals in trading sequences, which increases the likelihood of misclassification.

As the embedding dimension increases, Precision gradually improves, with the most notable gains observed between 64 and 128. This suggests that higher representation dimensions enhance the model's ability to capture features. The model becomes more effective in handling multidimensional order flows and microstructural characteristics. By expanding the embedding space, the model can better retain potential risk signals, thereby improving detection accuracy.

At an embedding dimension of 128, Precision reaches its highest value. This shows that a favorable balance between information richness and model complexity is achieved at this level. The model has sufficient expressive capacity without introducing excessive redundancy or noise, which makes its performance in risk detection more stable. This result confirms the sensitivity of performance to embedding dimensions and highlights the decisive role of proper parameter selection in improving accuracy.

When the embedding dimension continues to increase to 192 and above, Precision starts to decline. This indicates that overly large embedding spaces introduce redundant features and additional noise, and may also cause overfitting during training. In high-frequency trading environments, which are both high-dimensional and fast-changing, excessively large dimensions reduce detection efficiency and weaken the model's ability to focus on key risk signals. Therefore, a moderate embedding dimension is a critical factor for ensuring a balance between accuracy and efficiency.

This paper also gives the impact of the missing rate on the experimental results, and the experimental results are shown in Figure 4.

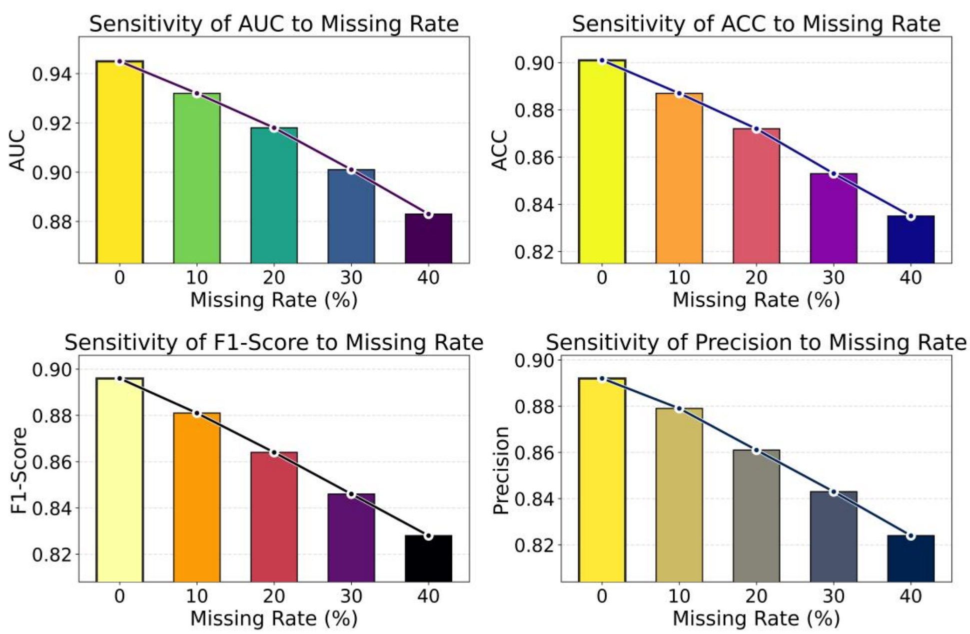

Figure 4. The impact of the missing rate on experimental results

From the experimental results, it can be observed that as the missing rate increases, all four metrics show a gradual decline. This indicates that data integrity plays a critical role in model performance for risk detection tasks in high-frequency trading. When data loss is minimal, the model can fully utilize multidimensional features for risk modeling, maintaining high accuracy and stability. However, as the missing rate rises, the loss of key information weakens feature representation, making it difficult for the model to accurately capture potential risk signals in the microstructure.

The decline in AUC and ACC is particularly notable. This suggests that missing data directly affects the model's ability to define classification boundaries and maintain accuracy. Once the missing rate exceeds a certain threshold, the model's capacity to distinguish between risk samples and normal samples weakens, leading to an overall drop in performance. This finding shows that in high-frequency environments, even minor data loss can cause systematic degradation in detection accuracy. It also highlights the need to minimize the impact of missing data during preprocessing.

The downward trend of F1-Score further reveals the effect of missing data on balanced performance. Since F1-Score combines Precision and Recall, increased missing data lowers both the recognition rate of anomalies and the reliability of predictions. This dual effect reduces the robustness of the model in complex market environments. It suggests that missing data weakens the ability of risk detection to capture true signals while suppressing noise.

The changes in Precision confirm this conclusion. As the missing rate increases, the model is more prone to misclassifying anomalies, reducing the reliability of detection results. In the context of high-frequency trading risk detection, this means the model may issue too many false alarms about market fluctuations, which reduces its practical value. Therefore, the experimental results indicate that improving data quality and completeness is the foundation for stable risk identification at the microstructural level. It is also a necessary condition for ensuring that risk detection frameworks function effectively in real markets.

This paper also gives the influence of noise intensity on the experimental results, and the experimental results are shown in Figure 5.

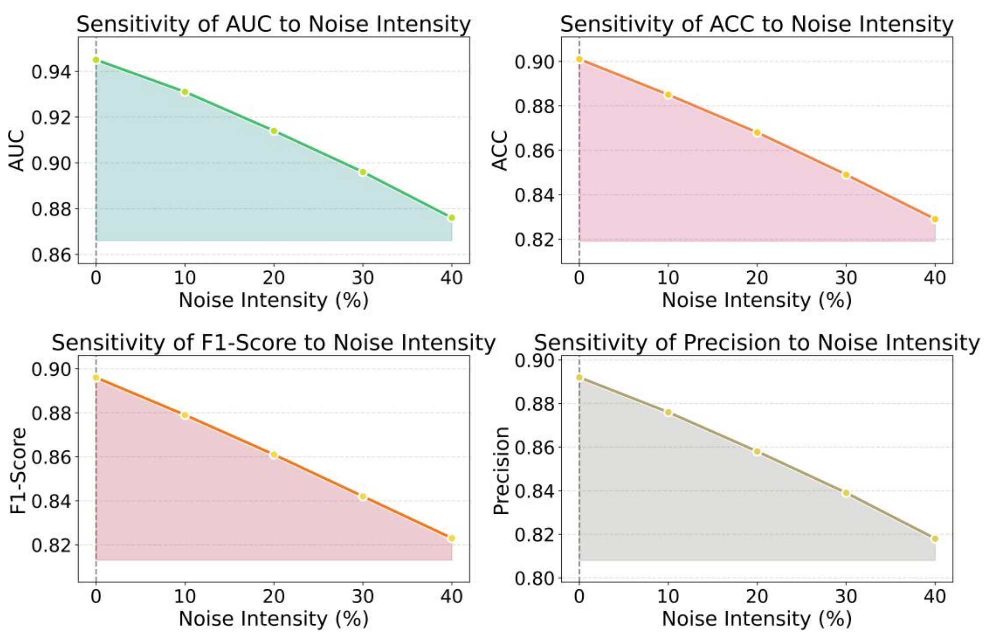

Figure 5. The influence of noise intensity on experimental results

From the experimental results, it can be seen that as noise intensity increases, the model's performance on the AUC metric continues to decline. This shows that noise directly weakens the model's ability to distinguish overall classification boundaries. At low noise levels, the model can stably identify risk signals. However, when noise accumulates to a certain degree, the feature differences between risk and non-risk samples are gradually obscured. As a result, the model's discriminative power in complex trading environments is significantly reduced.

The trend in the ACC metric further confirms this conclusion. As noise increases, overall prediction accuracy shows a clear decline. This indicates that data disturbances have a strong impact on the stability of the model. In high-frequency trading scenarios, this means that even if a model performs well under ideal conditions, once exposed to strong noise, it may produce large numbers of misjudgments in trading streams. This undermines both the reliability and the practicality of risk detection.

From the perspective of F1-Score, noise affects both recall and precision, creating a combined negative effect that lowers overall balanced performance. The continuous decline in F1-Score shows that when handling high-noise data, the model struggles to cover more true risk samples and, at the same time, generates more false alarms. This dual negative effect is critical in high-frequency financial contexts, where risk signals are often sparse and weak. Amplified noise directly reduces the model's ability to respond at crucial moments.

The changes in Precision highlight the impact of noise on the credibility of predictions. As noise levels rise, the reliability of risk identification continues to decrease, and the model is more likely to misclassify noise as abnormal events. In the context of high-frequency trading risk detection, this means that excessive noise produces too many false warnings, which reduces the value of detection systems in real markets. Therefore, the experimental results indicate that algorithm design must carefully account for noise robustness to ensure that models maintain high detection accuracy even when information quality is degraded.

5. Conclusion

This study conducts a systematic investigation of microstructural risk detection in high-frequency trading environments, covering both methodological modeling and sensitivity experiments. By modeling high-frequency features of trading sequences, the study demonstrates that capturing price dynamics and order flow characteristics at the micro level is crucial for effective risk detection in complex financial markets. The experimental results show that the proposed framework achieves superior performance across multiple metrics. It maintains stable detection under high-dimensional, nonlinear, and real-time conditions, thus offering new insights for financial risk identification.

In the comparative experiments, the study validates the differences between traditional methods and deep models under different conditions. The results indicate that algorithms based on attention mechanisms and dynamic modeling achieve significant improvements in both performance and robustness. This conclusion not only extends the theoretical boundaries of financial time series analysis but also provides practical references for market supervision and trading risk control. The study highlights the necessity of feature modeling and temporal dependency capture, emphasizing the value of microstructural modeling in enhancing risk recognition. It also lays the groundwork for further exploration of multi-source data and complex patterns in related fields.

In the sensitivity experiments, the study reveals the strong impact of factors such as hyperparameters, data completeness, and noise levels on model performance. Proper settings for sequence length and embedding dimension can balance information representation and computational complexity. Missing rates and noise intensity directly affect robustness and detection reliability. These findings demonstrate not only the superiority of the algorithm under ideal conditions but also the challenges it may face in real-world environments. They provide important experience and insights for the practical implementation of risk detection systems.

Overall, by constructing a microstructural risk detection framework and conducting comprehensive validation, this study offers solid theoretical and methodological support for financial risk identification in high-frequency trading environments. The findings enrich interdisciplinary research in financial technology and provide technical references and tools for regulators, trading system designers, and risk management departments. Against the backdrop of increasingly complex and fast-moving financial systems, the significance of this study lies in advancing risk detection methodologies toward greater refinement and intelligence, thereby contributing to market security and operational efficiency.

References

[1] J. Sirignano and R. Cont, "Universal Features of Price Formation in Financial Markets: Perspectives From Deep Learning," Quantitative Finance, vol. 19, no. 9, pp. 1449 – 1459, 2019.

[2] M. Goudarzi and F. Bazzana, "Identification of high-frequency trading: A machine learning approach," Research in International Business and Finance, vol. 66, Art. no. 102078, 2023.

[3] L. Lucchese, M. S. Pakkanen, and A. E. D. Veraart, "The short-term predictability of returns in order book markets: A deep learning perspective," International Journal of Forecasting, vol. 40, no. 4, pp. 1587-1621, 2024.

[4] A. Arroyo, A. Cartea, F. Moreno-Pino, et al., "Deep attentive survival analysis in limit order books: Estimating fill probabilities with convolutional-transformers," Quantitative Finance, vol. 24, no. 1, pp. 35-57, 2024.

[5] H. Hultin, H. Hult, A. Proutiere, et al., "A generative model of a limit order book using recurrent neural networks," Quantitative Finance, vol. 23, no. 6, pp. 931-958, 2023.

[6] Y. Wu, M. Mahfouz, D. Magazzeni, et al., "Towards robust representation of limit order books for deep learning models," arXiv preprint arXiv:2110.05479, 2021.

[7] R. James, H. Leung and A. Prokhorov, "A Machine Learning Attack on Illegal Trading," Journal of Banking & Finance, vol. 148, Art. no. 106735, 2023.

[8] G. Ibikunle, B. Moews, D. Muravyev, et al., "Data-driven measures of high-frequency trading," arXiv preprint arXiv:2405.08101, 2024.

[9] F. Mai, S. Tian, C. Lee and L. Ma, "Deep Learning Models for Bankruptcy Prediction Using Textual Disclosures," European Journal of Operational Research, vol. 274, no. 2, pp. 743–758, 2019.

[10] B. Dong, D. Zhang, and J. Xin, "Deep reinforcement learning for optimizing order book imbalance-based high-frequency trading strategies," Journal of Computing Innovations and Applications, vol. 2, no. 2, pp. 33-43, 2024.

[11] M. Nimalendran, K. Rzayev, and S. Sagade, "High-frequency trading in the stock market and the costs of options market making," Journal of Financial Economics, vol. 159, Art. no. 103900, 2024.

[12] B. Clapham, M. Haferkorn, and K. Zimmermann, "The impact of high-frequency trading on modern securities markets," Business & Information Systems Engineering, vol. 65, no. 1, pp. 7-24, 2022.

[13] S. Sun, "TrendFormer: Two-stage trend fusion for financial time series forecasting using transformers," Artificial, 2024.

[14] R. Yan and M. Guo, "Enterprise audit risk identification through temporal context and entity relationship anomaly modeling," 2024.

[15] Y. Nie, "Financial risk identification using unified multi-source feature learning and structural aggregation," 2024.

[16] Z. Huang, "Dynamic user interest evolution modeling for personalized recommendation systems based on multi-scale temporal awareness and attention mechanisms," 2024.

[17] Z. Pan, "Intelligent decision-making for advertising recommendation via data mining and graph neural networks," 2024.

[18] Y. Zhang and L. Steven, "Automatic orchestration of enterprise ETL processes via graph neural network and large language model collaborative reasoning," 2024.

[19] W. Li, S. S. Cheng, L. Mao, et al., "A benchmark study for limit order book (LOB) models and time series forecasting models on LOB data," 2024.

[20] W. Gu, Y. Zhong, S. Li, et al., "Predicting stock prices with FinBERT-LSTM: Integrating news sentiment analysis," in Proc. 2024 8th International Conference on Cloud and Big Data Computing, 2024, pp. 67-72.

[21] S. Lai, M. Wang, S. Zhao, et al., "Predicting high-frequency stock movement with differential transformer neural network," Electronics, vol. 12, Art. no. 2943, 2023.

[22] Y. Tang, S. Ding, and D. Zhang, "Integrating CNNs and transformers for mid-price prediction in high-frequency trading," in Proc. BenchCouncil International Symposium on Intelligent Computers, Algorithms, and Applications, Singapore: Springer Nature Singapore, 2024, pp. 61-77.