Transactions on Computational and Scientific Methods | Vo. 4, No. 01, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Financial Risk Identification Using Unified Multi-Source Feature Learning and Structural Aggregation

Yunfei Nie
Brandeis University, Waltham, USA
samnie519@gmail.com

Abstract: This paper proposes a supervised risk probability modeling framework based on multi-source feature fusion for financial risk identification tasks, enabling stable identification in business environments characterized by high noise, strong nonlinearity, and subject heterogeneity. The method takes multi-source data, such as transaction and identity data, as input, obtaining consistent numerical representations through a unified standardization and missing data correction process. A lightweight nonlinear encoder is then employed to learn risk-related latent space features. When structural information is available, neighborhood aggregation is introduced to inject local contextual relationships, thereby simultaneously characterizing individual attributes and interactive influences. Subsequently, the model maps the fused representation to risk logarithms and outputs risk probabilities via a Sigmoid algorithm, facilitating integration with threshold-based early warning and risk stratification strategies. During the training phase, weighted binary cross-entropy combined with a regularization term is used to adapt to class imbalance and differences in misjudgment costs, improving the robustness of the discrimination boundary. Comparative experiments on public datasets demonstrate that the proposed method outperforms existing representative models in accuracy, precision, recall, and overall metrics, validating the effectiveness and versatility of the framework in multi-source fusion and risk identification.

Keywords: Risk probability modeling; multi-source feature fusion; neighborhood aggregation; weighted cross-entropy

1. Introduction

Financial risk identification and early warning are core components in banking, securities, insurance, and regulatory scenarios. Their goal is to promptly identify potential risks and support sound decision-making in highly uncertain, nonlinear, and significantly heterogeneous business environments. With the rapid development of digital finance, multi-source data, including transaction behavior, account relationships, fund flows, and external information, continues to grow, leading to risk patterns characterized by cross-entity linkages, cross-temporal evolution, and cross-channel diffusion. Traditional risk management methods relying on rules and static indicators struggle to fully characterize complex interconnected structures and dynamic processes, making them prone to identification lags and misjudgments when facing new risk patterns, behavioral camouflage, and environmental disturbances, thus affecting the timeliness of risk management and system stability [1, 2].

Risk in real-world financial systems is not driven by a single variable but is a complex phenomenon formed by the coupling of multiple dimensions, encompassing both abnormal behavior at the individual level and synergistic effects and structural transmission at the group level. Risk signals are often implicit in high-dimensional feature spaces and interaction relationships, exhibiting problems such as weak separability, strong

noise, class imbalance, and sample distribution drift. Furthermore, risk events exhibit significant time-series dependence and phased characteristics. Macroeconomic policies, market sentiment, business strategies, and adversarial operations all alter risk generation mechanisms, necessitating models to maintain stable judgment capabilities and transferability under changing environments. This places higher demands on the expressive power, robustness, and generalization ability of risk identification methods [3].

Artificial intelligence-based risk discrimination methods offer a new technical path to address these challenges. Their key lies in learning complex patterns and underlying structures in a data-driven manner, and forming operable risk scores and discrimination results through multi-source information fusion and dynamic modeling [4]. Through stronger feature representation capabilities and nonlinear fitting capabilities, these methods can extract more risk-sensitive discrimination clues from high-dimensional data, reducing reliance on manual rules and subjective experience. Simultaneously, introducing structural and temporal information helps capture inter-subject relationships and risk transmission chains, thereby completing risk characterization in a representation space closer to real financial behavior and improving adaptability to hidden and emerging risk forms [5, 6].

In the context of high leverage, strong interconnectivity, and high-frequency operation in the financial system, constructing a highly reliable risk discrimination system is of great significance. On the one hand, accurate and robust risk identification capabilities help to expose potential default and fraud risks in advance, reducing institutional losses and improving operational security [7]. On the other hand, in response to regulatory and compliance risk monitoring needs, intelligent judgment capabilities can support more granular risk profiling and continuous monitoring, promoting a shift in risk management from post-event response to pre-event prevention. More importantly, given the ever-evolving financial business ecosystem, risk identification methods need to maintain stable output amidst complex noise and changing distributions, providing long-term support for the resilience and risk governance capabilities of the financial system.

2. Method

For the task of financial risk identification, the input samples consist of multi-source features, including numerical business indicators, behavioral statistical features, and optional structural correlation information. First, the original features are uniformly scaled, and missing data are processed to obtain a standardized vector $x_i \in R^d$. To reduce the impact of different scales on the discrimination boundary, mean-variance normalization is applied to each dimension, mapping the samples to a stable feature space as the basic representation for subsequent modeling. Figure 1 also shows the overall model architecture.

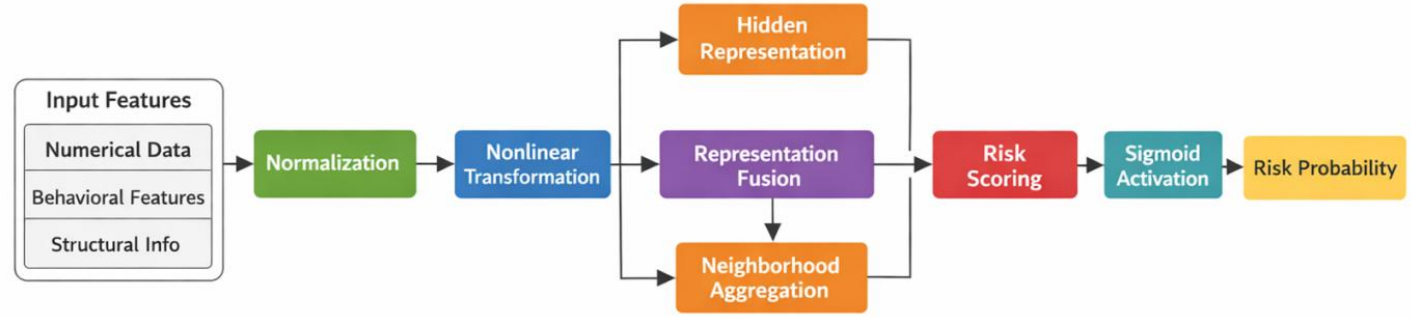

Figure 1. This paper outlines the proposed risk prediction architecture, which proceeds from multi-source feature input to probability output. The process performs feature normalization and non-linear encoding, selectively integrates neighborhood aggregation to obtain structural contextual information, and generates the final risk probability through risk scoring and a sigmoid activation function

This step ensures the numerical stability of the training process and facilitates consistent processing of new data in dynamic business environments, thereby improving the comparability and transferability of the model output. Standardization is defined as:

$$x_{i,k} = \frac{x_{i,k} - \mu_k}{\sigma_k + \varepsilon}$$

In the feature encoding stage, a lightweight nonlinear mapping is used to project the input representation onto the latent space, obtaining the intermediate representation h_i of the node or sample. This encoder uses linear transformations and element-wise activation functions to capture nonlinear relationships with low complexity while avoiding excessive parameter overhead. The encoding process is written as follows:

$$h_i = \phi(W\tilde{x}_i + b)$$

Where $\phi(\cdot)$ is the element-wise activation function, and W and b are learnable parameters. If the data contains inter-subject relationships or transactional relationships, a simple neighborhood aggregation can be performed in the latent space to inject the association structure information into the representation, so that the representation simultaneously includes individual attributes and local interaction effects. The aggregated representation of sample i is defined as:

$$g_i = \frac{1}{|N(i)| + 1}(h_i + \sum_{j \in N(i)} h_j)$$

Where $N(i)$ is the set of neighbors connected to i; in a scenario without structural information, $g_i = h_j$ can be obtained.

In the risk scoring layer, the fused representation is mapped to a scalar risk logarithm, and the risk probability $p_i \in (0,1)$ is obtained via Sigmoid, used for binary risk classification. This output form can be directly used for risk warning threshold strategies and subsequent risk stratification management, while maintaining interpretable probabilistic semantics. The risk probability is defined as:

$$p_i = \sigma(w^T g_i + c)$$

Where $\sigma(z) = \frac{1}{1+exp(-z)}$, w, and c are learnable parameters. If samples arrive in a time-series manner and short-term noise fluctuations need to be suppressed, a simple exponential smoothing of the probability sequence can be performed to obtain a more stable time-consistent output.

$$\tilde{p}_t = \alpha\tilde{p}_{t-1} + (1 - \alpha)p_t$$

Where $\alpha \in [0,1]$ is the smoothing coefficient and $\tilde{p}_t$ represents the smoothing probability at time t.

The objective function employs weighted binary cross-entropy to accommodate class imbalance and the difference in misclassification costs, and uses L_2 -regularization to suppress overfitting, ensuring stable discrimination across different business distributions. The sample loss for the label $y_i \in \{0,1\}$ is defined as:

$$L = -\sum_i \begin{matrix}(w_1 y_i log p_i + \\ w_0(1 - y_i)log(1 - p_i)) + \lambda||\theta||_2^2\end{matrix}$$

Where w_1 and w_0 are class weights, θ represents the set of all learnable parameters, and λ is the regularization coefficient. Through the above process from input normalization, nonlinear encoding, optional structure aggregation to probability output and robust optimization, a unified end-to-end risk discrimination method framework can be formed, supporting the stable risk identification needs under multi-source features and dynamic scenarios.

3. Experimental Results and Analysis

3.1 Dataset

This study utilizes the IEEE-CIS Fraud Detection open-source dataset for financial risk discrimination modeling. This dataset is designed for transaction fraud probability prediction tasks, with the core label "isFraud" and providing two tables that can be linked by TransactionID: the 'transaction' table contains multi-dimensional attribute features at the transaction level, and the 'identity' table contains identity-related information related to the device and network environment; however, not every transaction has corresponding identity information. This dataset is large in scale and rich in feature dimensions, covering multiple risk-related signals in typical online transaction scenarios, making it suitable as a general benchmark for financial risk discrimination.

In terms of data usage, the two tables are left-joined using TransactionID to construct a unified sample view. Missing identity fields are considered as one of the optional sources of structural information to maintain compatibility with different information completeness levels. The features as a whole include numerical transaction information and categorical context fields, directly supporting end-to-end risk probability modeling processes based on standardization and non-linear coding, while also providing a natural data entry point for subsequent context aggregation based on relationships or neighborhoods.

3.2 Dataset Preprocessing

(1) Data Integration and Field Alignment: The 'transaction' and 'identity' tables are left-joined by 'TransactionID' to form a unified sample view; fields used only for indexing or unsuitable for modeling are deleted, and field types are standardized, separating obvious string-type categorical fields from numeric fields to ensure consistency in subsequent coding and normalization processes.

(2) Missing Value Handling and Anomaly Correction: Numerical features are imputed using the median or grouped median to reduce the impact of extreme values; for categorical features, missing values are treated as independent categories and filled with placeholders; infinity and outlier values are truncated or set as missing before imputation, while missing value indicator features are recorded to retain the risk information carried by the missing value pattern. Figure 2 shows a comparison of the distribution of representative numerical features before and after preprocessing.

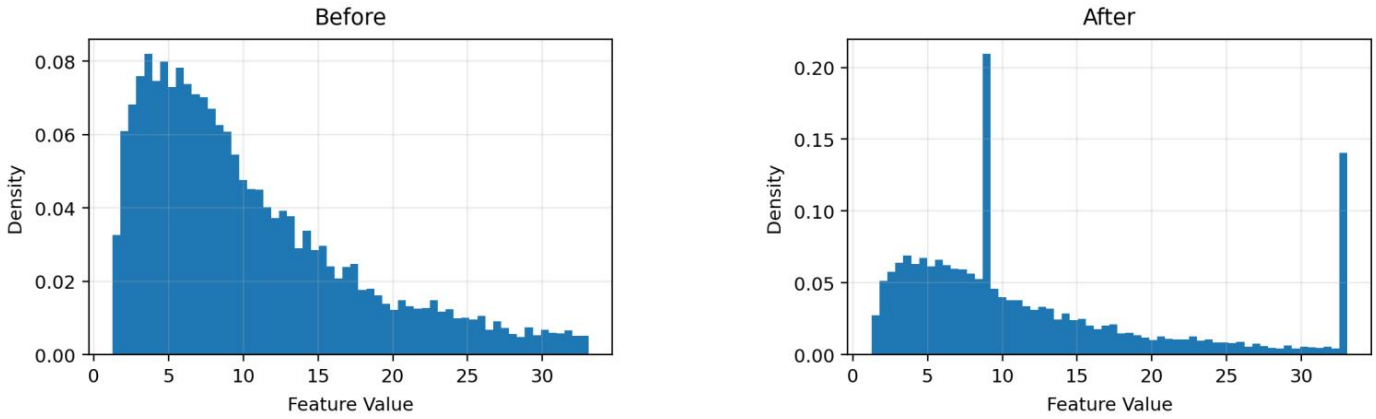

Figure 2. The distribution of representative numerical features before and after preprocessing is compared. After anomaly correction and missing value imputation based on the median, the distribution is more regular and stable

(3) Category Coding and Feature Construction: One-hot encoding is used for low-cardinality category features, while frequency encoding or target-independent hash encoding is used for high-cardinality category features to control dimensionality; logarithmic transformation is used for long-tailed numerical features such as

monetary amounts to alleviate skewness. Figure 3 shows a comparison of high cardinality features before and after frequency encoding and long-tail monetary features before and after logarithmic transformation.

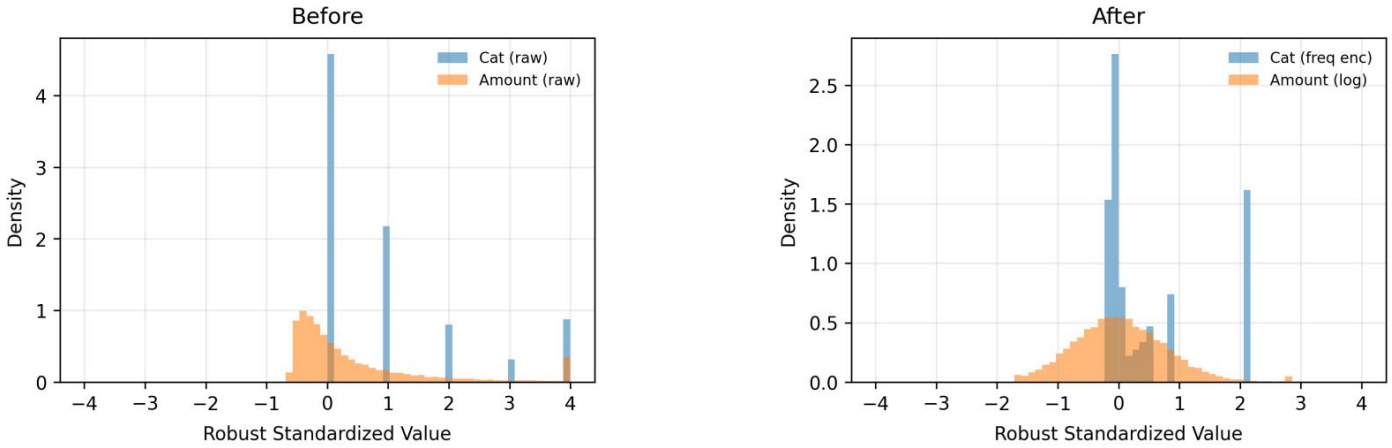

Figure 3. After frequency encoding of high cardinality features and logarithmic transformation of long-tail monetary features, the distribution of these features becomes more concentrated and regular, which facilitates subsequent stable modeling

(4) Feature Scaling and Data Partitioning: Mean and variance standardization or robust scaling is applied to numerical features to unify the dimensions; time-order or hierarchical partitioning strategies are used for training and testing to ensure that the label ratio is consistent with the business distribution, and a validation set is further partitioned within the training set; all normalization and encoding parameters are fitted only on the training set and applied to the validation and testing sets to avoid information leakage.

3.3 Experimental Results and Analysis

To fully demonstrate the methodological positioning and research context of the proposed method in the task of financial risk identification, representative research works in the same direction were selected for horizontal comparison, covering typical paradigms such as graph structure modeling, class imbalance learning, uncertainty characterization and tabular deep learning, to ensure that the comparison objects are comparable in terms of problem setting and data form. The experimental results are shown in Table 1.

Table 1. Experimental results compared with other models

Method	Accuracy	Precision	Recall	F1
Cheng et al. [8]	0.87	0.86	0.85	0.85
Dou et al. [9]	0.88	0.87	0.86	0.86
Liu et al. [10]	0.89	0.88	0.87	0.87
Tian et al. [11]	0.90	0.89	0.88	0.88
Habibpour et al. [12]	0.86	0.85	0.84	0.84
de la Bourdonnaye et al. [13]	0.88	0.86	0.87	0.86
Meng et al. [14]	0.91	0.90	0.89	0.89
Ours	0.94	0.93	0.92	0.92

The overall comparison reveals a clear performance gradient. Different models show largely consistent and synchronous changes across the four evaluation metrics, indicating a relatively stable trend in the trade-off between discriminative ability and error rate. Solutions based on graph structure modeling or designed for fraud-adversarial features typically perform well, achieving a better balance between precision and recall. Solutions that fail to adequately characterize class imbalance and complex interactions are relatively weaker, resulting in limited overall performance metrics.

Among the listed methods, the proposed method is optimal across all four metrics, demonstrating stronger overall discriminative ability and stability. This advantage is not only reflected in the improvement of individual metrics but also in the further increase in the overall metrics when precision and recall are improved simultaneously. This means that the risks of false positives and false negatives are more consistently controlled, and the model output better meets the robustness requirements of risk discrimination.

From the perspective of consistency, the proposed method leads across multiple metrics simultaneously, indicating that it more fully represents key risk signals and is less sensitive to noise and distribution fluctuations, thus maintaining a relatively stable decision boundary under different sample types. Compared to comparative methods that rely solely on a single representation or a single learning strategy, this method has advantages in information fusion and discriminant function construction, thus achieving superior overall performance.

The uncertainty smoothing coefficient is used to suppress short-term noise disturbances in the time dimension, thereby obtaining a more stable risk probability output. Since smoothing involves a trade-off between response speed and stability, different coefficient settings will change the model's sensitivity to sudden fluctuations and persistent patterns. To characterize the impact of this hyperparameter on the model's discriminative stability, a sensitivity analysis of the uncertainty smoothing coefficient on accuracy was conducted, and the experimental results are shown in Figure 4.

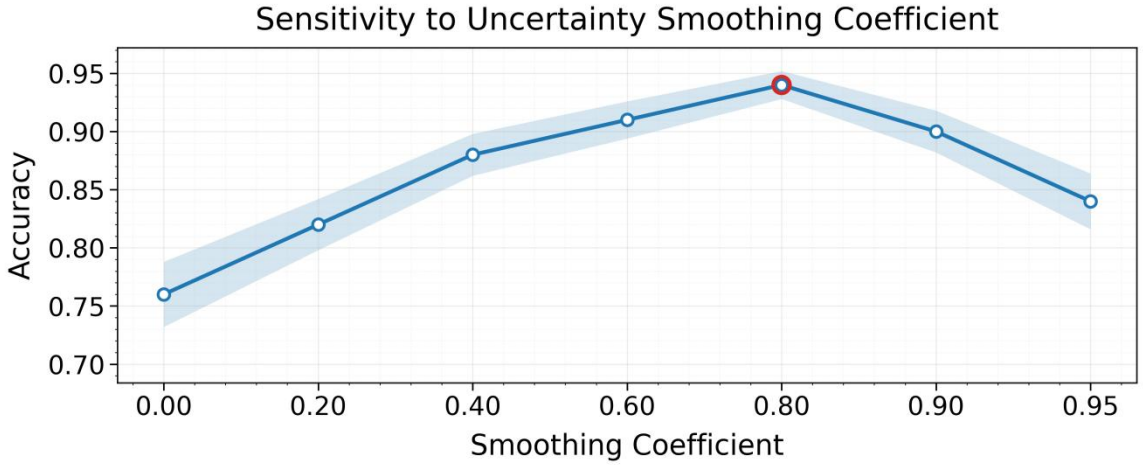

Figure 4. Experiment on the sensitivity of the uncertainty smoothing coefficient to accuracy

The curve shape shows that the uncertainty smoothing coefficient has a significant impact on the model's discriminative performance, exhibiting an overall trend of first increasing and then decreasing. A lower smoothing level is insufficient for noise suppression, making the output more susceptible to short-term disturbances. As smoothing intensifies, predictions become more stable, the discrimination boundaries become more consistent, and accuracy improves, reaching a better range. The change in the shaded band also reflects that under appropriate smoothing strength, the model's performance is more focused, and fluctuations are more controllable.

When the smoothing coefficient continues to increase, the curve declines, indicating that excessive smoothing introduces a hysteresis effect, weakening the response to changes in key signals, resulting in information loss and a decline in discriminative ability. This phenomenon suggests that a balance needs to be struck between stability and sensitivity in the smoothing term. Choosing a moderately high but not excessive coefficient can balance noise suppression and effective information preservation, providing a more robust parameter selection basis for practical deployment.

Noise injection intensity is used to characterize the robustness and stability of the model under perturbed environments. It simulates uncertainties in real-world business scenarios by controlling the amplitude of random perturbations on the input or representation side. Different intensities alter the smoothness of the feature

distribution and the consistency of local structures, thus affecting the stability of the discrimination boundary. To quantify the impact of this hyperparameter on model performance under perturbed scenarios, a sensitivity analysis of noise injection intensity on accuracy is conducted.

As shown in Figure 5, the model accuracy generally decreases with increasing noise injection intensity, and the decrease is relatively smooth, indicating that random perturbations continuously weaken the separability of input features and the consistency of local structure. Low-intensity noise has a relatively controllable impact on the discrimination boundary and can still retain the main risk signals; however, when the perturbation intensifies further, the feature distribution is more significantly flattened and aliased, making it more difficult for the model to stably capture effective patterns, thus exhibiting more significant performance degradation.

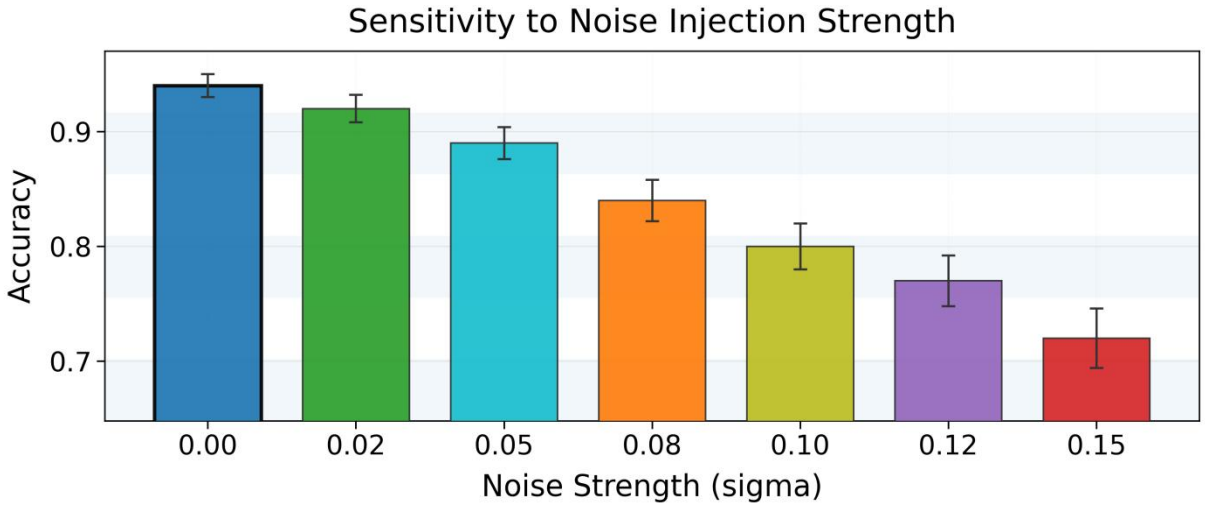

Figure 5. Sensitivity experiment of noise injection intensity to accuracy

The changes in the error bars indicate that the stronger the noise, the more easily the fluctuations in the results are amplified, reflecting an increased sensitivity of the model to random perturbations and a decrease in output stability. This phenomenon suggests that noise injection should be used as a robustness-enhancing adjustment term rather than being as strong as possible. A more moderate intensity is more beneficial in improving the ability to resist perturbations while avoiding excessive information loss, and provides a reasonable parameter selection basis for dealing with environmental noise and data perturbations in practical deployments.

4. limitation

This method belongs to a supervised risk discrimination framework, which is highly dependent on the quality and coverage of annotations. In real-world business scenarios, risk types are constantly evolving and exhibit conceptual drift. Historical annotations may not fully cover new fraud patterns or marginal risk forms, leading to potential blind spots in the model's identification when facing unseen distributions and novel adversarial strategies. Furthermore, class imbalance and sample selection bias amplify the instability of minority class learning, limiting the model's generalization ability in low-occurrence risk scenarios.

While the structural association and neighborhood aggregation introduced in the method can enhance contextual representation, their effectiveness depends on the availability and reliability of relational information. If the subject association graph has missing edges, noisy edges, or changes in connection mechanisms due to cross-domain migration, neighborhood aggregation may introduce bias propagation, weakening the robustness of the discrimination boundary. Simultaneously, data involving multi-source identity fields and device environment fields often suffer from missing information and differences in collection strategies. The encoding and construction strategies in the preprocessing stage may need recalibration under

different collection environments. This issue is summarized in Table 2 as the main source of environmental and data sensitivity.

From an engineering implementation perspective, while end-to-end probability output facilitates threshold-based early warning, high-risk decision-making scenarios still require more interpretable evidence chains and auditability support, such as key feature contributions, correlation paths, and uncertainty source decomposition. The current framework relies heavily on external analytical methods for explanation and has not yet formed a unified causal-level explanation structure within the model. To address these shortcomings, future work will focus on stronger out-of-distribution robustness, relational noise suppression, cross-environment adaptive calibration, and the generation of interpretable risk evidence. Specific limitations and potential improvement paths are shown in Table 2.

Table 2. Limitation analysis and potential improvement directions

Limitation	Impact on performance	Potential improvement directions
Dependence on labeled coverage in supervised learning	May lead to recognition blind spots for unseen risk patterns	Introduce semi-supervised and weakly supervised learning, combined with online updating and drift detection
Class imbalance and selection bias	Unstable learning for minority classes, making the false-positive/false-negative trade-off harder	Combine cost-sensitive learning with resampling, add hard example mining, and adaptive thresholding.
Missing or noisy relational information	Neighborhood aggregation may propagate bias and reduce robustness	Graph structure denoising and confidence-weighted aggregation, neighborhood filtering via consistency constraints
Cross-environment distribution shift	Encoding and statistical features require recalibration	Domain-adaptive calibration and feature alignment, environment-wise normalization, and robust encoding
Insufficient interpretability and auditability	Hard to form a traceable evidence chain to support decisions	Introduce explainability modules and evidence generation, output key features, and summaries of related paths
Privacy and compliance constraints	Multi-source identity fields may be restricted, reducing data availability	Federated learning and privacy-preserving modeling minimize feature sets and apply secure aggregation

5. Conclusion

It integrates standardized processing, nonlinear representation learning, optional aggregation of structural information, and probabilistic risk output end-to-end to adapt to the complex environment of real-world business, characterized by high noise, strong nonlinearity, and subject heterogeneity. The framework emphasizes robustness and transferability in its methodological design. Through concise yet effective modular combinations, the model can capture risk signals at the individual level and depict local interactive effects when relational and contextual information is available, thus forming a discriminative representation that more closely reflects the financial risk generation mechanism. Based on a unified risk probability output format, the method can naturally connect to downstream processes such as threshold-based early warning, risk stratification management, and continuous monitoring, providing operable technical support for risk governance.

At the application level, the proposed method provides a scalable risk identification solution for scenarios such as banking, payment, credit, internet finance, and regulatory technology. Facing diverse risk tasks such as transaction fraud, account anomalies, money laundering leads, and credit defaults, this framework supports cross-business reuse with consistent input representations and output semantics, reducing reliance on manual

rules and mitigating the problem of rapidly escalating rule maintenance costs with business expansion. Simultaneously, probabilistic outputs facilitate integration with business risk control strategies, supporting flexible threshold configuration and alarm grading under different risk appetites and resource constraints, thus ensuring both risk control strength and operational efficiency and user experience. For regulatory and compliance requirements, this method also provides a more granular modeling foundation for continuous monitoring and risk profiling, helping to shift risk governance from ex-post handling to ex-ante prevention and process control.

This research also has broader implications for intelligent upgrades in related fields. Financial risk inherently possesses strong correlation, dynamics, and adversarial characteristics; single-perspective features or static rules are unlikely to be effective in the long term. By fusing multi-source information within a unified framework and maintaining optional access methods for structured information, the method can better adapt to differences in data completeness and system construction stages among different institutions, providing a feasible path for the gradual upgrade of large-scale risk control systems. Furthermore, the framework's modular design facilitates integration with existing data governance, feature platforms, and online service architectures, supporting a smooth migration from offline scoring to online real-time risk control and enhancing the automation and intelligence of the overall risk management chain.

Looking to the future, it remains necessary to continuously improve the framework in a direction that more closely aligns with real-world business constraints. On one hand, it is necessary to further enhance its adaptability to new risk patterns and distribution changes, forming a stronger cross-time and cross-scenario generalization mechanism to address the ongoing challenges brought about by iterative risk strategies and evolving adversarial behaviors. On the other hand, more efficient security modeling paradigms should be explored under privacy protection and compliance constraints to improve multi-institutional collaborative risk governance capabilities and reduce data sharing costs. Simultaneously, interpretability and auditability will become crucial requirements in high-risk decision-making scenarios. Future work should focus on enhancing the usability and verifiability of model outputs in areas such as generating traceable evidence chains, attributing key factors, and characterizing risk propagation paths, thereby further enhancing the model's practical application value and social impact within financial institutions and regulatory systems.

References

[1] Y. Liu, Z. Sun and W. Zhang, "Improving Fraud Detection via Hierarchical Attention-Based Graph Neural Network," Journal of Information Security and Applications, vol. 72, Art. no. 103399, 2023.

[2] K. Li, T. Yang, M. Zhou et al., "SEFraud: Graph-Based Self-Explainable Fraud Detection via Interpretative Mask Learning," Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 5329-5338, 2024.

[3] W. Hyun, I. Lee and B. Suh, "LEX-GNN: Label-Exploring Graph Neural Network for Accurate Fraud Detection," Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, pp. 3802-3806, 2024.

[4] G. Tong and J. Shen, "Financial Transaction Fraud Detector Based on Imbalance Learning and Graph Neural Network," Applied Soft Computing, vol. 149, Art. no. 110984, 2023.

[5] L. Wang, Z. Cheng, M. Yang et al., "GEM-GNN: Group Enhanced Multi-Relation Graph Neural Networks for Fraud Detection," Proceedings of the International Conference on Advanced Data Mining and Applications, pp. 275-290, 2024.

[6] Y. Zhao, "Cross-Timescale Transformer With One-Dimensional Convolution for Integrated Financial Risk Anomaly Detection and Discrimination," 2024.

[7] A. Roy, J. Sun, R. Mahoney, L. Alonzi, S. Adams and P. Beling, "Deep Learning Detecting Fraud in Credit Card Transactions," Proceedings of the Systems and Information Engineering Design Symposium (SIEDS), pp. 129-134, 2018.

[8] D. Cheng, X. Wang, Y. Zhang et al., "Graph Neural Network for Fraud Detection via Spatial-Temporal Attention," IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 8, pp. 3800-3813, 2020.

[9] Y. Dou, Z. Liu, L. Sun et al., "Enhancing Graph Neural Network-Based Fraud Detectors Against Camouflaged Fraudsters," Proceedings of the 29th ACM International Conference on Information and Knowledge Management, pp. 315-324, 2020.

[10] Y. Liu, X. Ao, Z. Qin et al., "Pick and Choose: A GNN-Based Imbalanced Learning Approach for Fraud Detection," Proceedings of the Web Conference 2021, pp. 3168-3177, 2021.

[11] Y. Tian, G. Liu, J. Wang et al., "Transaction Fraud Detection via an Adaptive Graph Neural Network," arXiv preprint arXiv:2307.05633, 2023.

[12] M. Habibpour, H. Gharoun, M. Mehdipour et al., "Uncertainty-Aware Credit Card Fraud Detection Using Deep Learning," Engineering Applications of Artificial Intelligence, vol. 123, Art. no. 106248, 2023.

[13] F. De La Bourdonnaye and F. Daniel, "Evaluating Categorical Encoding Methods on a Real Credit Card Fraud Detection Database," arXiv preprint arXiv:2112.12024, 2021.

[14] C. C. Meng, K. M. Lim, C. P. Lee et al., "Credit Card Fraud Detection Using TabNet," Proceedings of the 2023 11th International Conference on Information and Communication Technology (ICoICT), pp. 394-399, 2023.