

Article

Not peer-reviewed version

Anomaly Ranking for Enterprise Finance Using Latent Structural Deviations and Reconstruction Consistency

[Hejing Chen](#), Ruibo Wu, Chen Chen, [Hanrui Feng](#), Yunfei Nie, [Yixuan Lu](#)*

Posted Date: 3 March 2026

doi: 10.20944/preprints202603.0197.v1

Keywords: representation learning; outlier scoring; reconstruction consistency; robustness assessment



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.

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.

Article

Anomaly Ranking for Enterprise Finance Using Latent Structural Deviations and Reconstruction Consistency

Hejing Chen ¹, Ruibo Wu ², Chen Chen ³, Hanrui Feng ⁴, Yunfei Nie ⁵ and Yixuan Lu ^{6,*}

¹ Walsh College, Troy, USA

² University of California, San Diego, La Jolla, USA

³ Vanderbilt University, Nashville, USA

⁴ University of Chicago, Chicago, USA

⁵ Brandeis University, Waltham, USA

⁶ University of Sofia, Palo Alto, USA

* Correspondence: yixuan.lu17@gmail.com

Abstract

This paper targets key challenges in enterprise financial anomaly monitoring, including coupled multi-source heterogeneous features, concealed anomaly signals, and cost-sensitive alerting. It proposes a unified monitoring framework based on representation-driven data modeling. The approach first normalizes and encodes transaction and identity-related fields within a unified feature space. A learnable representation mapping function is then used to obtain low-dimensional latent representations. A reconstruction consistency mechanism is introduced to suppress noise while preserving essential behavioral information. In the latent space, normal behavior structures are statistically characterized. Structural anomaly signals are derived by measuring the deviation between representation centers and sample representations. These signals are combined with reconstruction errors from the original inputs to form a unified anomaly score. This enables continuous ranking of samples and threshold-based anomaly decisions. To ensure controllability in training and deployment, the framework is systematically analyzed from the perspectives of hyperparameter sensitivity, including learning rate, optimizer choice, and network depth, as well as environmental sensitivity under input noise injection. The analysis delineates the impact boundaries of key configurations on overall detection performance and stability. The results demonstrate that the framework maintains consistent risk identification characteristics under different configurations and perturbations. It provides sortable and configurable anomaly signals for enterprise risk alerts, audit sampling, and compliance inspection.

Keywords: representation learning; outlier scoring; reconstruction consistency; robustness assessment

CCS CONCEPTS: Computing methodologies-Machine learning-Machine learning approaches

1. Introduction

Against the backdrop of the digital economy and the continuous deepening of enterprise informatization, corporate financial activities are generating data at an unprecedented speed and scale [1]. Higher transaction frequency, increasing business complexity, and cross-platform and cross-entity data flows have exposed the limitations of traditional rule-based, experience-driven, and post hoc auditing approaches. These methods are often delayed and lack sufficient coverage. Financial anomalies rarely appear as isolated points. They are embedded in high-dimensional, multi-source, and dynamically evolving data structures. They often manifest as subtle but persistent pattern shifts. This reality makes it critical to systematically distinguish normal and abnormal behaviors from

complex financial data. The problem has become a core challenge in enterprise risk management and governance[2].

At the same time, corporate financial data exhibit strong structural and relational properties. Accounting records, account behaviors, business interactions, and temporal evolution are not independent. They are tightly coupled through implicit business logic and organizational relationships. Methods that rely on single-indicator thresholds or static feature statistics struggle to capture such cross-dimensional and cross-level dependencies. This limitation often leads to both false alarms and missed detections[3]. As anomalous behaviors become more concealed and fragmented, financial risks no longer present as extreme values. They are more often reflected in shifts of overall representations or disruptions of local relations. This trend places higher demands on modeling capability for anomaly monitoring.

In this context, representation-based data modeling offers a new perspective for corporate financial anomaly monitoring. By mapping raw financial data into a unified, low-noise, and discriminative representation space, it becomes possible to describe the global structure and dynamic characteristics of financial behaviors at an abstract level. This approach moves beyond detecting abnormal fluctuations of individual variables. It focuses on consistency and stability in multidimensional feature combinations, behavior patterns, and evolution trajectories within the representation space. As a result, it is better suited to identify hidden anomaly signals under complex conditions. Anomalies are no longer treated as deviations from a mean. They are understood as deviations from normal structural patterns[4].

From the perspective of corporate governance and risk control, building a financial anomaly monitoring framework based on representation modeling has clear practical value. On the one hand, such a framework enables earlier awareness of financial risks[5]. Continuous monitoring of structural changes in the representation space can provide earlier and more fine-grained risk signals for management decisions. On the other hand, a unified representation space supports the integration of risk information across business units and data sources. It helps reduce information silos and improve consistency and coordination in enterprise-wide risk management[6]. This representation-centered paradigm supports a shift from passive compliance-oriented management to proactive risk prevention.

Moreover, under stricter regulatory requirements and rising business uncertainty, financial anomaly monitoring affects not only internal efficiency but also compliance and long-term stability. A representation-based anomaly monitoring framework can maintain generalization ability while adapting to ongoing changes in business structures and external environments[7]. It provides flexible technical support for financial risk identification in complex settings. By modeling financial behaviors at the representation level, such a framework can serve as a key link between data assets, risk understanding, and governance decisions. It holds long-term significance for improving financial transparency and strengthening enterprise risk resilience.

2. Proposed Framework

2.1. Overall Modeling Framework Overview

This paper proposes a corporate financial anomaly monitoring framework based on representation-driven data modeling. The method aims to characterize the structural regularities of enterprise financial behaviors in a unified latent space and to identify anomalies through distribution modeling and structural stability analysis. The original multi-source heterogeneous financial data are first normalized and encoded into a unified feature space, enabling cross-source semantic alignment and structural coupling preservation.

In constructing the representation mapping mechanism, this study draws methodological inspiration from the embedding-based structural reasoning paradigm proposed by Ying et al [8]. Their work fundamentally models complex relational dependencies by embedding structured entities into a continuous representation space for causal reasoning over knowledge graphs. Building

upon this structural abstraction principle, our study adopts embedding transformation strategies to unify transaction records, identity attributes, and contextual financial indicators into a shared latent space. We apply learnable nonlinear mappings to encode heterogeneous attributes, incorporate relational coupling constraints to preserve implicit business logic, and leverage structural embedding mechanisms to ensure that latent representations reflect cross-feature interactions rather than isolated indicators. Unlike causal intervention modeling, we extend the representation paradigm toward structural stability modeling for anomaly detection.

To suppress noise while preserving essential behavioral information, a reconstruction consistency mechanism is introduced. The residual-regulated learning strategy proposed by Ou et al. [9] fundamentally controls predictive stability by modeling residual structures and mitigating non-stationary drift through second-order differencing. We adopt the principle of residual regulation and apply it to representation learning by incorporating reconstruction residual monitoring as a structural stability constraint. The reconstruction module is not used merely for dimensionality compression; instead, we leverage residual signals to distinguish structural deviations from random fluctuations, build upon residual stabilization strategies to enhance robustness under dynamic financial environments, and extend their non-stationary modeling intuition to representation-level behavioral modeling.

In the latent space, normal financial behaviors are characterized through statistical structural modeling. The adaptive anomaly detection framework proposed by Ou et al. [10] introduces continual learning with dynamic distribution monitoring to address evolving data distributions. Their method fundamentally tracks distributional shifts and updates detection boundaries adaptively. We adopt this dynamic distribution monitoring mechanism and incorporate it into latent structural characterization. Specifically, we apply adaptive center estimation and leverage continual monitoring strategies to track representation drift. By building upon dynamic distribution modeling principles, the framework maintains consistent anomaly scoring boundaries under environmental perturbations and evolving business conditions. We extend continual adaptation from time-series forecasting contexts to representation-based structural anomaly identification. Assuming the original corporate financial data consists of multi-source features, it can be represented as:

$$X = \{x_1, \dots, x_n\} \quad (1)$$

Each x_n represents a financial feature vector aggregated within a specific time window or business unit. The overall framework maps the original features to a low-dimensional latent space using a representation mapping function.

$$z_i = f_\theta(x_i) \quad (2)$$

where f_θ is a learnable representation function, and $z_i \in R^k$ and $k < d$. This latent representation space is considered the core carrier for characterizing the normal financial behavior structure of an enterprise, and subsequent anomaly monitoring is carried out in this space.

This article presents the overall model architecture diagram, as shown in Figure 1.

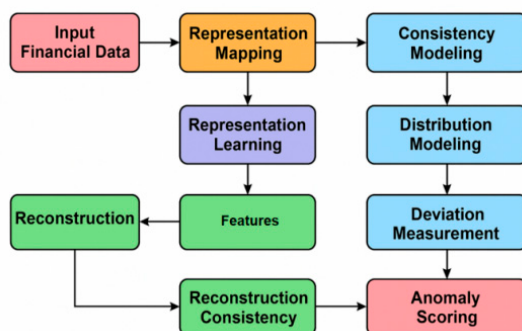


Figure 1. Overall model architecture diagram.

2.2. Financial Behavioral Representation Learning Modeling

To ensure that the representation space can stably depict normal financial behavior patterns, this paper adopts a representation learning approach based on reconstruction consistency to constrain the latent representation. Suppose there exists a reconstruction mapping function $g_\phi(\cdot)$, then the reconstruction process can be represented as:

$$x_i = g_\phi(z_i) \quad (3)$$

The corresponding reconstruction error is defined as:

$$L_{rec} = \frac{1}{n} \sum_{i=1}^n \|x_i - \hat{x}_i\|_2^2 \quad (4)$$

This constraint prompts representation z_i to retain key financial information while filtering out redundant noise, thereby forming a compact and comparable behavioral representation. Through this mechanism, the model can learn the implicit structure that reflects the company's routine financial operating logic, laying the foundation for subsequent anomaly detection.

2.3. Representing Spatial Structure Consistency Modeling

Within the latent representation space, this study assumes that statistically stable financial behaviors form a structurally coherent distribution rather than scattered individual points. Building upon the structured embedding and generative abstraction mechanism proposed by Long et al. [11], which fundamentally models semantically consistent representations through knowledge-guided embedding generation, we adopt semantic-constrained representation regularization to ensure that normal samples aggregate around a structurally meaningful center. At the same time, leveraging the dynamic structural dependency modeling strategy introduced by Gan et al. [12], which captures evolving relational interactions via causal graph mechanisms, we incorporate adaptive structural awareness into distribution estimation to prevent center distortion under temporal and cross-entity variations. In addition, drawing on the attention-driven feature weighting framework proposed by Wang et al. [13], which fundamentally enhances anomaly discrimination through adaptive importance allocation, we apply attention-aware refinement prior to center estimation to suppress noise and strengthen compactness of normal clusters. Accordingly, the representation distribution is modeled as a structurally constrained and adaptively maintained central manifold, and the central vector of this distribution is defined as:

$$\mu = \frac{1}{n} \sum_{i=1}^n z_i \quad (5)$$

Based on this center, an offset metric is defined to characterize the distance between a single financial sample and the overall structure:

$$s_i = \|z_i - \mu\|_2 \quad (6)$$

This metric reflects the degree to which current financial behavior deviates from the overall structure in the representation space, providing an intuitive and stable quantitative basis for anomaly detection. By performing structural modeling at the representation level rather than the original feature level, the interference of high-dimensional noise and local fluctuations on anomaly judgment can be effectively reduced.

2.4. Financial Anomaly Scoring and Judgment Mechanism

After completing representation learning and structural modeling, anomaly identification is reformulated as a representation-shift-driven scoring process in which deviations are quantified through integrated structural and reconstruction signals. Building upon the causal modeling and consistency-aware learning framework proposed by Li et al.[14], which fundamentally mitigates correlation bias by enforcing consistency constraints under causal regularization, we adopt consistency-aware integration to jointly evaluate reconstruction stability and latent structural deviation, apply bias-mitigated scoring principles to prevent spurious correlations from dominating

anomaly assessment, and incorporate causal-inspired regularization to enhance robustness of the unified score. Leveraging the Deep Q-learning-based intelligent scheduling strategy introduced by Gao et al.[15], which fundamentally optimizes decision policies through reward-driven adaptive control in heterogeneous environments, we incorporate adaptive weighting mechanisms into the scoring formulation, apply dynamic adjustment strategies to balance reconstruction error and representation shift, and build upon reinforcement-based optimization principles to enhance configurability under varying deployment conditions. Drawing on the transformer-based heterogeneous data modeling approach proposed by Xie and Chang, [16] which fundamentally captures cross-feature dependencies through attention-driven contextual aggregation, we adopt cross-dimensional interaction modeling to ensure that anomaly scoring reflects global representation coherence rather than isolated feature deviations. Furthermore, inspired by the adaptive risk control mechanism in the multi-agent reinforcement learning framework proposed by Ying et al. [17], which fundamentally coordinates agents under dynamic risk constraints, we incorporate adaptive risk sensitivity into the scoring boundary design, leverage dynamic adjustment principles to stabilize anomaly thresholds, and extend risk-aware optimization strategies toward unified anomaly quantification. Accordingly, considering both reconstruction consistency and representation structural shift, a unified anomaly scoring function is defined as follows:

$$A_i = \alpha \|x_i - x_i\| + (1 - \alpha) \|z_i - \mu\|_2 \quad (7)$$

Here, α is the balance coefficient, used to adjust the relative weights of original behavioral consistency and representational structural stability in anomaly detection. This scoring mechanism avoids over-reliance on a single indicator, allowing anomaly identification to be based on the overall behavioral representation, thus better reflecting the actual characteristic that corporate financial risk often manifests as structural shifts rather than extreme numerical changes.

3. Experimental Analysis

3.1. Dataset

This study adopts the open source IEEE CIS Fraud Detection dataset as a unified benchmark for enterprise financial anomaly monitoring. The dataset is designed for real-world transaction risk control scenarios. It provides structured records at the transaction level. Anomalies can be directly defined as high-risk or fraudulent transactions, which represent a typical form of financial abnormal events. The task setting closely matches the objective of identifying a small number of abnormal behaviors from massive transaction streams. It is therefore suitable for evaluating the effectiveness and practical feasibility of representation space modeling for anomaly detection.

The dataset consists of two main components. The transaction table contains business-related features such as transaction time, transaction amount, product category, payment card attributes, address information, and distance-related variables. It also includes a large number of anonymized statistical and behavioral features. The identity table provides device type, device information, and multiple anonymized identity-related variables. These two tables can be linked through transaction identifiers. This linkage forms a typical multi-source input structure. It supports the modeling setting of mapping multiple fields and multiple source data into a unified latent representation space.

From the perspective of anomaly monitoring requirements, the dataset exhibits properties of multidimensional heterogeneity, relational constructability, and temporal aggregatability. Multiple feature types can be jointly encoded into transaction-level representation vectors. Fields such as card identifiers, email domains, and device information naturally induce implicit links across transactions. This property facilitates the construction of representation structures that better reflect real enterprise risk propagation patterns. In addition, transaction timestamps enable window-based aggregation to form stable behavioral segment representations. This aligns with the anomaly scoring strategy based on representation consistency and structural shift.

Overall, the dataset provides a reproducible and publicly verifiable open benchmark for enterprise financial anomaly monitoring. It enables systematic evaluation without reliance on proprietary accounting systems. This makes it suitable for validating representation-based monitoring frameworks in a transparent and repeatable manner.

3.2. Experimental Setup

All training and inference were conducted in a single machine and single-GPU setting. The hardware included one NVIDIA GeForce RTX 4090 GPU with 24GB memory, an AMD Ryzen 9 7950X processor, 128GB system memory, and a 2TB NVMe SSD. The software environment consisted of Ubuntu 22.04 LTS, Python 3.10.13, PyTorch 2.1.2, CUDA 12.1, and cuDNN 8.9. A fixed random seed of 42 was used in all runs. The number of data loading workers was set to 8. Automatic mixed precision training was enabled. Float16 was used during training, while float32 was adopted during inference to ensure stable scoring behavior.

The hyperparameter configuration is summarized as follows. Input preprocessing included zero-filling for missing values, standardization for continuous features, and logarithmic scaling for amount-related variables. Categorical features were encoded using learnable embeddings with dimension 16. The representation mapping network was implemented as a three-layer multilayer perceptron with a hidden width of 256 and a representation dimension of 128. ReLU was used as the activation function, and the dropout rate was set to 0.10. The reconstruction decoder adopted a symmetric two-layer structure with a hidden width of 256. The AdamW optimizer was applied with a learning rate of 0.0005 and weight decay of 0.01. The first and second moment coefficients were set to 0.9 and 0.999. The model was trained for 100 epochs with a batch size of 2048. Gradient clipping was applied with a threshold of 1.0. A cosine annealing learning rate schedule was used, with linear warm-up during the first five epochs. The anomaly threshold was defined as the 99.5th percentile of the score distribution on the training set. This threshold was used to convert continuous anomaly scores into binary anomaly labels.

3.3. Experimental Results

This article first presents the results of the comparative experiments, as shown in Table 1.

Table 1. Comparative experimental results.

Method	Acc	Precision	Recall	F1-Score
CoDetect [18]	0.87	0.86	0.81	0.83
LAMDA [19]	0.89	0.88	0.83	0.85
Tcnattention-rag [20]	0.90	0.89	0.85	0.87
Weirdnodes [21]	0.91	0.90	0.86	0.88
Ours	0.93	0.92	0.89	0.91

A global comparison reveals a clear and stable performance hierarchy among the baseline methods. This pattern indicates that the task is highly sensitive to representation quality and to the definition of anomaly decision boundaries. The proposed approach achieves the best results on Acc, Precision, Recall, and F1. In particular, the F1 score reaches 0.91. This result reflects a better balance between false alarm control and missed anomaly reduction. Such a balance aligns with the core requirements of enterprise financial anomaly monitoring. In high noise settings with multiple source features and low anomaly ratios, it is critical to reduce unnecessary disruption caused by false alarms while avoiding losses from undetected high-risk behaviors.

The improvement in Precision suggests that the method provides clearer separation between normal and abnormal samples in the anomaly scoring space. Anomaly decisions are concentrated in more reliable risk regions. For enterprise financial monitoring, higher Precision implies more actionable alerts. Audit and investigation resources are less likely to be wasted on low-value warnings. The proposed method reaches a Precision of 0.92. This further reduces false positives

compared to other approaches. The result indicates that the representation space captures normal financial behavior in a more compact manner. As a consequence, deviations can be identified more consistently.

At the same time, the gain in Recall indicates more comprehensive anomaly coverage. The model is able to capture hidden or weak signal anomalies. Financial anomalies often emerge as gradual structural shifts rather than abrupt extreme values. Models that rely only on local features or short-term patterns tend to miss such cases. The Recall of the proposed method reaches 0.89. This shows that anomaly detection is not limited to obvious outliers. It remains sensitive to marginal anomalies along potential risk chains. This behavior is consistent with the objective of a framework based on representation consistency and structural shift. It helps expand risk coverage under complex business relations and multidimensional behavior signals.

Overall, the improvement in F1 demonstrates a better trade-off in enterprise financial anomaly monitoring. Alert credibility is preserved while anomaly capture ability is enhanced. Compared with the incremental gains of baseline methods, the advantage of the proposed approach stems from a shift in modeling paradigm. Anomaly detection moves from raw feature thresholds and shallow pattern matching to measuring stability in latent representation structures. This leads to more robust decision boundaries under multi-source, multi-dimensional, and dynamic financial data distributions. Such characteristics are better suited for deployment in enterprise monitoring processes that are constrained by risk control objectives and limited resources. They provide more consistent signals for continuous monitoring and audit decision-making.

The learning rate determines the step size of parameter updates and is a key factor affecting the stability of representation learning and the convergence behavior of anomaly scoring boundaries. Different learning rates alter the model's trade-off between noise suppression and pattern fitting, thus affecting its sensitivity to identifying corporate financial anomalies. To verify the robustness and controllability of the proposed framework to changes in the learning rate, multiple learning rates were set, and the response trends of core indicators with the learning rate were observed. The experimental results are shown in Figure 2.

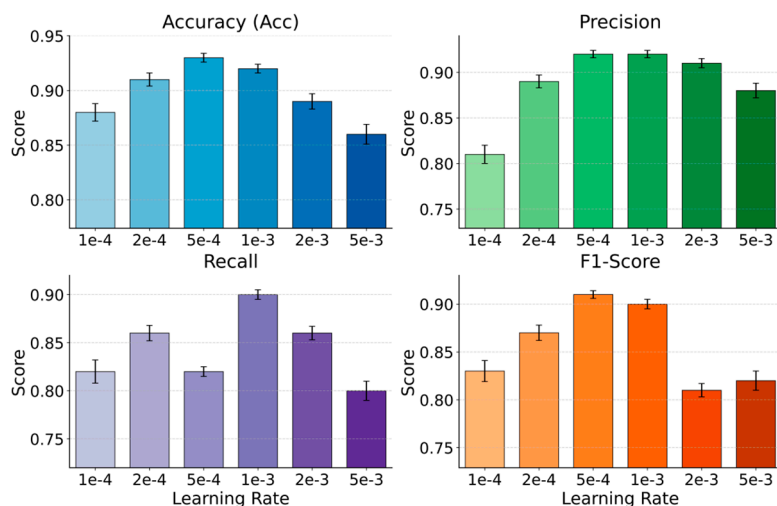


Figure 2. The impact of the learning rate on experimental results.

Across the four subplots, the overall shapes reveal a clear interval effect of the learning rate on representation learning effectiveness. When the learning rate is too low, parameter updates are limited. The model remains close to its initial representation structure. It struggles to absorb critical differences from multiple source financial behaviors. As a result, overall scores stay low and improve slowly. As the learning rate increases to a moderate range, all metrics rise together. This indicates that the latent representation space begins to form clearer normal structures and anomaly shift boundaries. The anomaly scoring function becomes more confident in separating risk samples.

From the perspective of accuracy, the bars associated with a moderate learning rate are clearly higher. This suggests a more stable global decision boundary and a more complete separation between normal and abnormal samples. For enterprise financial anomaly monitoring, this means the system is less likely to suffer from widespread misclassification during continuous operation. This is especially important when transaction distributions exhibit intraday fluctuations or business structure changes. An appropriate update step allows the representation space to remain adaptive while avoiding frequent boundary drift that would trigger unstable alerts. When the learning rate increases further, accuracy declines. This implies that overly aggressive updates disrupt the previously compact normal structure and reduce overall consistency.

For precision and recall, their optimal learning rate ranges do not fully coincide. Precision reaches a higher level around the moderate learning rate. This indicates that alerts are concentrated in more reliable risk regions. It reduces cases where normal transactions are incorrectly flagged as anomalies. This directly affects whether audit resources are consumed by low-value alerts. Recall becomes more prominent at a slightly higher but still moderate learning rate. This reflects a stronger tendency to capture marginal and weak signal anomalies, which expands risk coverage. When the learning rate continues to increase, recall drops sharply. This suggests instability in the representation space, where true structural shifts are masked by noisy updates.

The trend of the F1 score further confirms the typical trade-off in enterprise financial anomaly monitoring. False alarms must be controlled, while missed detections must be avoided. F1 reaches its peak in the moderate learning rate interval. This indicates that reconstruction consistency and structural shift measurement produce a more coordinated anomaly score. A better balance is achieved between alert reliability and anomaly coverage. For a framework centered on representation-based data modeling, this observation highlights the learning rate as a key control factor. It strongly influences representation geometry and threshold stability. Selecting a moderate learning rate is more conducive to producing controlled, reproducible, and stable outputs for real-world financial monitoring processes.

This paper also presents the impact of the optimizer on the experimental results, as shown in Table 2.

Table 2. The impact of optimizers on experimental results.

Optimizers	Acc	Precision	Recall	F1-Score
AdaGrad	0.88	0.86	0.82	0.84
Adam	0.91	0.90	0.86	0.88
SGD	0.89	0.87	0.84	0.85
AdamW	0.93	0.92	0.89	0.91

A comparison of different optimizers indicates that optimization strategies strongly influence representation learning and the stability of anomaly decision boundaries, with AdamW achieving the best performance across all metrics. Its adaptive learning rates and explicit regularization help handle heterogeneous feature scales and suppress overfitting, leading to more reliable risk representations and stable alerts. Compared with Adam, AdamW provides balanced improvements in Precision and Recall, reducing both false alarms and missed anomalies, which is critical for enterprise financial monitoring. In contrast, AdaGrad performs less effectively, likely due to overly rapid learning rate decay that limits adaptation under dynamic data distributions, while SGD shows moderate performance because of its lack of adaptivity and sensitivity to manual tuning in heterogeneous feature spaces. The F1-score ranking further confirms that AdamW offers the best trade-off between reliability and anomaly coverage, making it suitable as a default optimization strategy. In addition, the number of network layers directly affects representation capacity and structural abstraction, influencing the separation between normal patterns and abnormal shifts, and the impact of model depth on performance is analyzed in Figure 3.

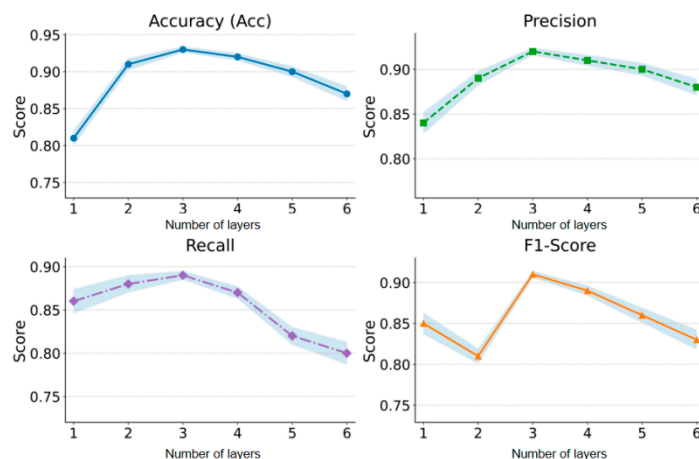


Figure 3. The impact of network layer number on experimental results.

The curve patterns indicate that network depth has a clear impact on the geometry of the latent representation space. An optimal complexity range can be observed. With a small number of layers, the expressive capacity of the mapping is limited. The model cannot fully absorb relational information across multiple source financial features. Normal structures are therefore not compact. Anomaly shifts are more easily obscured by mixed signals. As depth increases to a moderate level, the model learns more stable behavior representations. Structural shift signals used for anomaly scoring become clearer. Overall identification performance improves.

From the accuracy curve, a peak appears around a moderate depth and then gradually declines. This shows that deeper networks do not necessarily yield better decision boundaries. In enterprise financial anomaly monitoring, overly deep models are more likely to fit short-term noise and incidental patterns as if they were generalizable structures. Normal clusters in the representation space become more dispersed. Stable separation of anomaly shifts is weakened. This behavior is consistent with real business data, where feature dimensionality is high but effective signals are sparse. Excessive model capacity tends to introduce unnecessary complexity in decision boundaries.

The precision curve rises first, then reaches a plateau, and finally declines at larger depths. This implies that an appropriate depth allows alerts to concentrate in more reliable risk regions. False alarms and audit resource waste are reduced. When depth increases further, precision decreases. Boundary constraints for normal samples become weaker. Normal structures are overpartitioned. More normal transactions fall into high-risk regions. For a framework centered on representation consistency and structural shift, this suggests that encoder depth must match the noise level of the data. Otherwise, representation consistency is damaged, and alert credibility declines.

The trends of recall and F1 more clearly reflect the trade-off between anomaly coverage and alert reliability. Recall reaches a high point at moderate depth and then drops sharply. This indicates that overly deep networks may oversmooth anomaly features or be disrupted by noise fitting. Marginal and weak signal anomalies are missed. The F1 curve forms a peak near the moderate depth and declines when the network is too shallow or too deep. This shows that this depth range maintains both false alarm control and anomaly capture ability. It better satisfies enterprise requirements for stability and controllability. It also supports the conclusion that the proposed method more easily forms separable and interpretable representation structures under moderate model complexity.

The intensity of input noise injection can perturb the statistical structure and representation space consistency of financial features, thereby affecting the stability and environmental robustness of the anomaly scoring mechanism in this paper. The experimental results are shown in Figure 4.

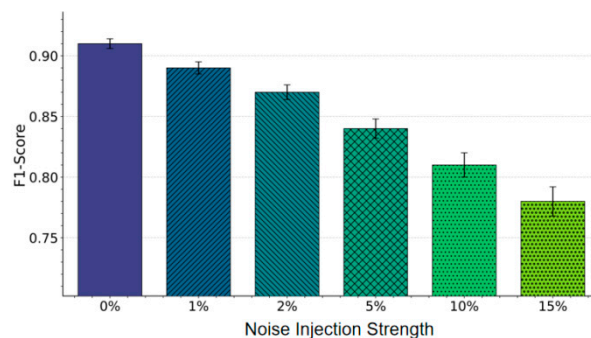


Figure 4. Environmental sensitivity analysis experiment of input noise injection intensity on F1.

The bar chart shows that as the intensity of input noise injection increases, the F1 score exhibits a consistent downward trend. This indicates that environmental disturbance directly weakens the overall anomaly detection capability. In enterprise financial anomaly monitoring, such noise can correspond to data collection errors, inconsistent field reporting, cross-system synchronization bias, and statistical perturbations introduced by anonymization. Increased noise blurs the previously clear boundaries of normal structures in the representation space. As a result, the responsiveness of anomaly scoring to structural shifts is reduced. In the low noise range, the F1 score remains at a relatively high level. This suggests that the representation learning mechanism has a certain degree of robustness to mild perturbations. Stable representations of normal behavior can still be maintained. From a business perspective, this means that when data quality fluctuates slightly, alert outputs do not immediately become unreliable and risk signals remain usable. Under a framework centered on representation consistency and structural shift, mild noise is not sufficient to disrupt reconstruction consistency or the central structure of representations. The relative geometry between normal and abnormal samples is largely preserved.

When noise intensity reaches a moderate level, the decline in F1 accelerates. This reflects a significant loss in separability between normal and anomalous samples in the representation space. Noise at this level not only perturbs features but also alters similarity structures among samples. The model struggles to maintain compact normal clusters. As a consequence, the score distributions of anomalies and normal samples begin to overlap. Once overlap increases, the system faces simultaneous risks of rising false positives and rising false negatives. The former consumes audit resources, while the latter conceals true abnormal behaviors. This directly targets the most sensitive cost and risk trade-off in enterprise financial monitoring. In the high noise range, the F1 score continues to decline and approaches a low level. This indicates that the tolerance limit of the model to strong disturbances has been exceeded. Structural stability required for representation modeling can no longer be preserved. For practical deployment, this result highlights the need to strengthen data pipeline quality control and consistency checks. Measures include missing value handling, outlier truncation, field standardization, and cross-source alignment. These steps reduce structural drift caused by noise injection. The results also show that representation-based anomaly monitoring frameworks have quantifiable robustness boundaries under environmental changes. The sensitivity curve can serve as an operational reference for alert credibility. It can guide decisions on when to trigger data remediation or strategy updates.

4. Discussion

The representation-based enterprise financial anomaly monitoring framework proposed in this study is designed to address the most common and challenging risk patterns in digitalized enterprise operations. Abnormal behaviors no longer appear as extreme values in a single field. They are embedded in structural shifts caused by multi-source data coupling, cross-entity relations, and temporal evolution. For enterprises, the value of anomaly monitoring lies not only in identifying

individual high-risk transactions. It also lies in the early detection of abnormal concentration regions and potential risk chains within massive daily business flows. Such signals can support financial shared service centers, internal control and compliance units, and audit teams with actionable risk clues. By placing monitoring in a unified representation space, weak anomalies scattered across transaction amounts, device information, identity attributes, and behavioral statistics can be integrated into consistent risk representations. This reduces the fragmented maintenance cost of rule systems when migrating across business lines. It also better aligns with enterprise needs for sustainable risk control systems. From a deployment perspective, the main constraints of financial anomaly monitoring systems often come from business processes and resource allocation rather than from theoretical model accuracy. In practice, excessive false positives trigger large volumes of manual review and business intervention. This directly leads to operational friction and degraded customer experience. Excessive false negatives may result in financial losses, regulatory penalties, and reputational damage. Enterprises, therefore, require a controllable monitoring mechanism. Alert intensity should be stably produced under data quality variation and business fluctuations. The system should also support integration with existing workflows. Different response strategies can be applied through risk tiering. High-risk samples can enter rapid freezing or secondary verification. Medium risk samples can be routed to manual sampling review. Low-risk samples can be retained for trace monitoring only. A representation modeling framework is more suitable in this respect. It compresses complex inputs into sortable anomaly scores. Thresholds can be adjusted to match the risk preferences of different departments. This enables dynamic balance between audit sampling budgets and risk coverage.

In addition, enterprise financial data inevitably contains noise and drift. These include cross-system definition differences, missing or delayed fields, business rule changes, and distribution shifts caused by seasonal activities. A key insight from the discussion is that environmental sensitivity analysis is not only an academic robustness test. It can also be transformed into an operational quality monitoring tool. It can characterize performance boundaries under different noise levels. It can guide data governance and model update schedules. For example, when data quality indicators show abnormal fluctuations, more conservative threshold strategies or additional consistency checks can be applied first. Model retraining or feature remapping can be performed when necessary. By linking representation space stability with business data quality management, anomaly monitoring systems can better integrate into enterprise risk management loops. They can support long-term goals in continuous compliance, fraud prevention, and financial resilience.

5. Conclusions

This study addresses key challenges in enterprise financial anomaly monitoring, including multi-source heterogeneity, strong structural dependency, concealed anomalies, and high cost sensitivity. It constructs a unified framework centered on representation-based data modeling. Anomaly detection is advanced from static judgments based on single-point thresholds and experience-driven rules to a systematic monitoring paradigm focused on the stability of latent representation structures. The framework builds unified representations of financial behaviors. It jointly considers consistency constraints and structural shift measures in the representation space. Anomalies are therefore not treated as extreme fluctuations in individual fields. They are characterized as disruptions and deviations from normal behavioral structures. This perspective is closer to real enterprise risk patterns. For practical business processes, the approach also supports continuous risk scores for tiered response strategies. It provides sortable and configurable decision bases for audit sampling, risk control interception, and compliance review. Manual review burden is reduced. Resource efficiency in risk handling is improved.

From an application impact perspective, the contribution extends beyond improved anomaly detection performance. It offers a scalable and transferable technical path for enterprise-level risk control systems. Enterprises often face inconsistent definitions, missing data, and distribution drift when integrating across business lines, regions, or systems. Traditional rule-based systems require

constant maintenance and struggle to maintain consistency. A representation-centered monitoring framework compresses multiple source fields and complex relations into a unified latent space. Risk identification relies more on structural signals rather than isolated features. This makes reuse across business scenarios easier. The paradigm also aligns naturally with existing enterprise data governance, internal control systems, and audit workflows. It can serve as a core risk engine for continuous monitoring. The focus shifts from post-event accountability to early warning and process-level control. It directly supports fraud prevention, fund security, financial compliance, and operational resilience.

Future work can focus on stronger adaptation to real-world scenarios and on interpretable governance. On the one hand, online updating and drift adaptive mechanisms can be introduced. The representation space can remain stable and controllable under business rule changes, macro environment fluctuations, or emerging anomaly strategies. Sensitivity curves can be linked with data quality monitoring to form an actionable operational loop. On the other hand, the structured interpretability of anomaly scores can be enhanced. Key driving factors, relational chains, and shift paths of high-risk samples can be output in auditable forms. This improves traceability and trust in compliance review and cross-department collaboration. As enterprise demand for privacy protection and cross-organization cooperation increases, future research can also explore joint representation learning under privacy constraints and cross-institution risk sharing. This would allow the framework to better support fraud prevention and collaborative risk defense in multi-party ecosystems. It can further expand its impact in financial technology, supply chain finance, enterprise digital internal control, and regulatory technology domains.

References

1. W. Hilal, S. A. Gadsden and J. Yawney, "Financial fraud: A review of anomaly detection techniques and recent advances," *Expert Systems with Applications*, vol. 193, p. 116429, 2022.
2. M. Ahmed, A. N. Mahmood and M. R. Islam, "A survey of anomaly detection techniques in financial domain," *Future Generation Computer Systems*, vol. 55, pp. 278-288, 2016.
3. M. Ahmed, N. Choudhury and S. Uddin, "Anomaly detection on big data in financial markets," *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2017.
4. A. Anandkrishnan et al., "Anomaly detection in finance: Editors' introduction," *KDD 2017 Workshop on Anomaly Detection in Finance*, 2018.
5. S. Crépey et al., "Anomaly detection in financial time series by principal component analysis and neural networks," *Algorithms*, vol. 15, no. 10, p. 385, 2022.
6. A. Elliott et al., "Anomaly detection in networks with application to financial transaction networks," *arXiv preprint arXiv:1901.00402*, 2019.
7. T. Chen and C. Tsourakakis, "Antibenford subgraphs: Unsupervised anomaly detection in financial networks," *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.
8. R. Ying, Q. Liu, Y. Wang and Y. Xiao, "AI-based causal reasoning over knowledge graphs for data-driven and intervention-oriented enterprise performance analysis," 2025.
9. Y. Ou, S. Huang, R. Yan, K. Zhou, Y. Shu and Y. Huang, "A residual-regulated machine learning method for non-stationary time series forecasting using second-order differencing," 2025.
10. Y. Ou, S. Huang, F. Wang, K. Zhou and Y. Shu, "Adaptive anomaly detection for non-stationary time-series: A continual learning framework with dynamic distribution monitoring," 2025.
11. S. Long, K. Cao, X. Liang, Y. Zheng, Y. Yi and R. Zhou, "Knowledge graph-driven generative framework for interpretable financial fraud detection," 2025.
12. Q. Gan, R. Ying, D. Li, Y. Wang, Q. Liu and J. Li, "Dynamic spatiotemporal causal graph neural networks for corporate revenue forecasting," 2025.
13. H. Wang, C. Nie and C. Chiang, "Attention-driven deep learning framework for intelligent anomaly detection in ETL processes," 2025.

14. S. Li, Y. Wang, Y. Xing and M. Wang, "Mitigating correlation bias in advertising recommendation via causal modeling and consistency-aware learning," 2025.
15. K. Gao, Y. Hu, C. Nie and W. Li, "Deep Q-learning-based intelligent scheduling for ETL optimization in heterogeneous data environments," arXiv preprint arXiv:2512.13060, 2025.
16. A. Xie and W. C. Chang, "Deep learning approach for clinical risk identification using transformer modeling of heterogeneous EHR data," arXiv preprint arXiv:2511.04158, 2025.
17. R. Ying, J. Lyu, J. Li, C. Nie and C. Chiang, "Dynamic portfolio optimization with data-aware multi-agent reinforcement learning and adaptive risk control," 2025.
18. D. Huang et al., "CoDetect: Financial fraud detection with anomaly feature detection," IEEE Access, vol. 6, pp. 19161-19174, 2018.
19. S. Zhang, Z. Feng and B. Dong, "LAMDA: Low-latency anomaly detection architecture for real-time cross-market financial decision support," Academia Nexus Journal, vol. 3, no. 2, 2024.
20. Y. Liu et al., "TCNAttention-RAG: Stock prediction and fraud detection framework based on financial report analysis," 2025.
21. S. Vilella et al., "WeirdNodes: Centrality based anomaly detection on temporal networks for the anti-financial crime domain," Applied Network Science, vol. 10, no. 1, pp. 1-29, 2025

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.