Transactions on Computational and Scientific Methods | Vol. 5, No. 6, 2025
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Deep Latent Representation and State Recursion for Enterprise Audit Risk Detection

Jiaming Chen[1]*
[1]Northeastern University, Boston, USA
*Corresponding author: Jiaming Chen, klauschen9@gmail.com

Abstract: In highly complex and data-intensive corporate operating environments, audit risk increasingly exhibits hidden structural patterns, temporal evolution, and multi-actor interactions. Traditional audit analysis methods that rely on static indicators and manual rules struggle to capture how risk forms and accumulates over time. To address this challenge, this study proposes a dynamic enterprise audit risk detection approach based on deep representation learning. The method performs unified modeling of multi-source corporate behavior data and maps enterprise operating states into a latent representation space. Risk states are continuously updated and characterized along the temporal dimension. Deep nonlinear mappings are used to automatically extract high-level semantic features. These features describe latent structural relationships and dynamic changes in corporate transaction behaviors. A state recursion mechanism is introduced to model the evolution of audit risk in a systematic manner. This avoids information fragmentation caused by single-point judgments. In addition, a unified risk scoring mechanism is constructed to express risk states in a continuous and comparable form. This enhances discrimination consistency and stability in complex audit scenarios. Comparative experimental results show that the proposed method outperforms representative existing approaches in overall identification performance and risk differentiation capability. The findings validate the effectiveness of combining deep representations with dynamic modeling for enterprise audit risk identification. This work provides a systematic solution for shifting audit risk analysis from static assessment to dynamic perception. It also offers a valuable reference for data-driven intelligent auditing research.

Keywords: Audit risk modeling; deep representation learning; dynamic risk detection; enterprise transaction analysis

1. Introduction

As business operations become increasingly complex and digitalized, the scale of internal and external transactions continues to expand. Business structures, capital flows, and interorganizational relationships exhibit strong nonlinearity and dynamic evolution[1]. In this context, audit risk no longer arises solely from anomalies in individual financial indicators. It is often embedded in complex behavioral patterns that span multiple entities, time periods, and business processes. Traditional audit risk identification methods that rely on rule-based constraints and static indicators struggle to capture implicit structural dependencies and evolutionary trends in corporate activities. As a result, risk exposure is often detected with significant delay, which fails to meet current regulatory and practical demands for proactive and continuous auditing[2].

At the same time, the data landscape of audit objects is undergoing a profound transformation. In addition to structured financial statements, large volumes of transaction records, contract texts, operational logs, and multi-source behavioral data are continuously accumulated. These data form a high-dimensional, heterogeneous, and temporally correlated space. They are characterized by large scale, high noise, unstable distributions, and diverse patterns. Under such conditions, analysis methods based on manual feature

engineering and linear assumptions are difficult to maintain stable performance. How to automatically extract deep representations that reflect the true operational state of firms and their potential risk characteristics has become a central challenge in audit risk research[3].

Advances in deep representation learning provide a new theoretical and technical foundation for audit risk modeling. Through multi-layer nonlinear mappings, models can learn high-level abstract features directly from raw data without explicit rule design[4]. They can capture latent structures and internal logic underlying corporate behaviors. This data-driven representation paradigm helps overcome the reliance of traditional audit methods on expert experience and fixed indicator systems. It enables a shift from surface-level anomaly detection toward a deeper understanding of behavioral patterns. This offers a more expressive modeling approach for complex audit scenarios.

However, audit risk is inherently dynamic in nature. Corporate strategies, transaction networks, and financial structures continuously evolve in response to changes in the external environment. Risk often emerges through gradual accumulation, structural shifts, or sudden disruptions. Static representations tend to overlook the evolutionary processes through which risks are formed[5,6]. They struggle to reflect changes in risk states over time. Therefore, integrating deep representation learning with temporal dynamic modeling to achieve continuous and dynamic characterization of audit risk is essential. This integration is critical for improving the timeliness and stability of risk identification.

Against this background, developing dynamic enterprise audit risk detection methods based on deep representation learning is of both theoretical and practical significance. Such research promotes a transition from static analysis to dynamic perception in audit risk studies[7]. It also provides more forward-looking technical support for audit practice. By enabling a systematic understanding of complex corporate behaviors and enhancing insight into the evolution of potential risks, this line of work can support regulatory decision-making, audit planning, and risk prevention with greater precision and continuity. It ultimately contributes to improved corporate governance and the advancement of intelligent auditing.

2. Proposed Framework

2.1 Overall framework modeling

This paper proposes a dynamic detection method for enterprise audit risk based on deep representation learning. The core idea is to map the multi-source operational behaviors of an enterprise at different times into a unified latent representation space, and to characterize the continuous evolution of risk status over time. Let the original observation data of the enterprise at time t be represented as a vector $x_t \in R^d$, which includes information such as financial indicators, transaction statistics, and behavioral characteristics. Through a nonlinear mapping function $f(\bullet)$, the original input is projected onto a low-dimensional latent space to obtain the deep representation of the enterprise at that time.

$$h_t = f(x_t)$$

This representation is considered an abstract depiction of the overall operating status of an enterprise, providing a unified state description basis for subsequent risk modeling and dynamic updates. This paper presents the overall model architecture, as shown in Figure 1.

2.2 Deep Representation Learning Module

In the representation learning phase, the model progressively extracts high-order semantic features through multi-layer mapping to reduce noise interference and enhance the ability to express key risk patterns. Assuming the representation learning module consists of layer L, the hidden representation of layer l is defined as follows:

$$z_t^{(l)} = \sigma(W^{(l)} z_t^{(l-1)} + b^{(l)})$$

Where σ are nonlinear activation functions. The final output firm state representation can be written as:

$$h_t = z_t^{(L)}$$

This process enables the model to automatically learn the nonlinear relationships between different features, thereby forming a more robust and discriminative potential representation of business operations.

Streaming time-series anomaly detection further motivates learning latent features that remain responsive as observations accumulate, rather than optimizing a static embedding for a fixed audit snapshot [8]. Hierarchical encoding with selective attention provides a complementary principle for preserving salient evidence from long transaction records and supporting documents while suppressing repetitive or low-value content in the representation space [9].

2.3 Dynamic modeling of risk status

Considering the time-varying nature of audit risk, this paper introduces a state recursion mechanism in the latent representation space to dynamically characterize the risk. Let the potential risk state of each enterprise at time t be s_t, then its update process is defined as:

$$s_t = g(s_{t-1}, h_t)$$

Where $g(\cdot)$ is the state transition function, used to integrate historical risk information with the current operating status. Furthermore, to characterize the magnitude of risk changes, the risk increment between adjacent time points can be defined as:

$$\Delta s_t = s_t - s_{t-1}$$

This dynamic modeling approach eliminates the reliance on single-point judgments for risk identification, instead reflecting the accumulation, mitigation, or sudden occurrence of risks through continuous state changes.

The state update is also designed to distinguish genuine risk evolution from incidental co-movement across corporate activities. Causal-aware anomaly detection offers a useful basis for tracing the evolution of operational states through causally informative transitions [10], while spatio-temporal graph forecasting supports representing how resource and behavioral signals propagate among dependent entities over successive observation periods [11].

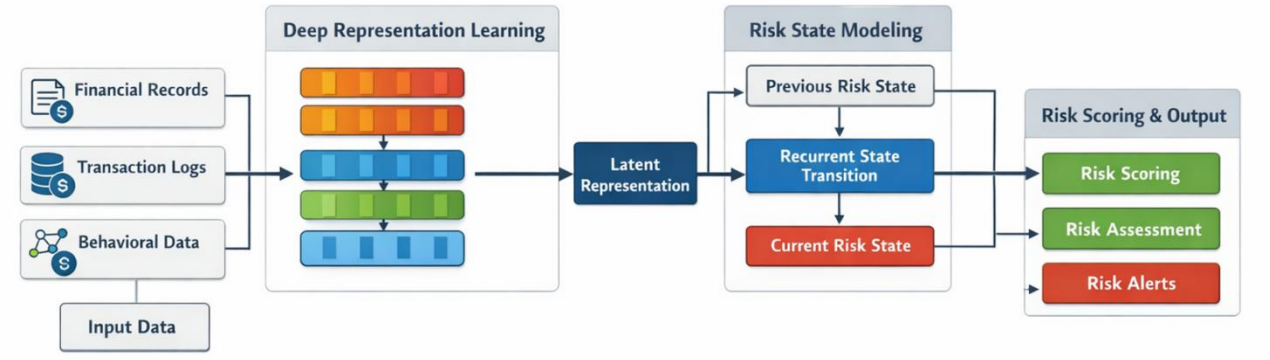


Figure 1. Overall model architecture diagram

2.4 Audit risk scoring and assessment

Based on the dynamic risk status, the model outputs an interpretable audit risk score to reflect the company's relative risk level at different times. Let the risk scoring function be $r(\bullet)$, then the company's risk value at time t is defined as:

$$R_t = r(s_t)$$

To ensure the stability and comparability of the scores, further normalization can be performed:

$$R_t = \frac{R_t - min(R)}{max(R) - min(R)}$$

Through the above modeling process, audit risk is characterized as a quantitative indicator that changes continuously over time, thus providing a consistent and dynamic risk representation basis for subsequent risk monitoring, early warning, and decision support.

The score can further support risk-responsive allocation of limited audit attention. Hierarchical reinforcement learning for microservice resource orchestration and adaptive scheduling under dynamic competition provide a decision-layer perspective for prioritizing investigations as risk states change [12], [13]. Interpretable feedback design in vision-language modeling also motivates presenting score changes with evidence-linked explanations, so that auditors can inspect the behavioral cues that influence each assessment [14].

3. Experimental Analysis

3.1 Dataset

This study uses a well-defined open source dataset, namely the AAER Dataset, also referred to as the Accounting and Auditing Enforcement Releases Dataset. The dataset is constructed from accounting and auditing enforcement releases issued by regulatory authorities. It is organized into structured samples that support the identification of financial misstatements and disclosure violations. The dataset contains a large number of enforcement documents and corresponding records of corporate misstatement events. It therefore provides a direct, authoritative, and traceable source of risk labels for research on enterprise audit risk.

In terms of data content, the AAER Dataset covers a long time span and systematically consolidates enforcement releases with firm-level misstatement events. This design facilitates the construction of a risk sample repository centered on individual firms. Each event is typically associated with a clearly specified violation type, the relevant time period, and affected financial statement elements. This structure allows data organization around the evolution of risk from occurrence to exposure and resolution. As a result, the dataset is naturally aligned with the problem setting of dynamic detection.

Within the task definition of this study, the AAER Dataset is used to construct dynamic supervised signals for enterprise audit risk. Firms are treated as indexing units, and their operational observations at different time points are mapped into a unified representation space. Misstatement events recorded in AAER serve as anchor points for risk states. They are used to characterize temporal trends and stage transitions in risk evolution. In this way, the dataset supports time series modeling of enterprise risk while satisfying open source and reproducibility requirements. It provides a reliable data foundation for dynamic audit risk detection under a deep representation learning framework.

3.2 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	AUC
FraudGNN-RL[15]	65.42	66.03	64.87	67.15

FraudGT[16]	67.58	68.11	66.92	69.34
Finstack-net[17]	68.73	69.24	67.85	70.12
Starlit[18]	69.05	69.87	68.43	71.08
FFDM-GNN[19]	70.21	71.03	69.56	72.44
Ours	**73.84**	**74.62**	**72.91**	**76.03**

From the overall trend, the performance of the comparative methods shows a stepwise improvement pattern. This indicates that audit risk identification in enterprise transaction networks benefits from enhanced graph structural modeling capability. Earlier approaches rely more on local neighborhood aggregation or relatively coarse relationship representations. As a result, they tend to mix true risk patterns with normal operational fluctuations when facing complex related party transactions, circular trades, and multi-actor collaborative behaviors. This leads to blurred decision boundaries. As the ability to model structural dependencies improves, the separability of risk samples in the representation space increases. This observation is consistent with the practical characteristic that audit risks are often embedded in multi-hop relationships and concealed paths.

More importantly, Ours maintains a leading position across all four evaluation metrics. This indicates that the method does not excel in only a single metric but achieves a more balanced overall discriminative capability. For dynamic audit risk detection, such a balance means that the model can remain stable under stricter thresholds while also reducing false alarms under more relaxed thresholds. This helps avoid unnecessary consumption of audit resources. The advantage at the AUC level is particularly noteworthy. It suggests that the model provides more stable ranking performance across different risk intensities and threshold ranges. This aligns well with audit practice, where high-risk firms or transactions must be prioritized from large populations.

The consistency across metrics further shows that the superior performance is not achieved by sacrificing recall to gain precision. It is also not obtained by increasing recall at the cost of excessive false positives. Instead, the method improves both the capture of risk samples and the discrimination of non-risk samples. This behavior matches the structural nature of enterprise audit risk. Risks often emerge as abnormal interaction patterns within relational networks. Effective detection requires uncovering hidden risk chains while avoiding misclassification of normal business collaboration. In this sense, the model learns a more generalizable structured risk representation rather than memorizing surface features or isolated anomaly signals.

In line with the theme of this paper, the results demonstrate that a structure modeling framework centered on deep representation learning can more effectively absorb dynamic semantics from enterprise transaction relationships. It can encode both the topology of interactions and the evolving behavioral patterns embedded in transaction sequences. This capability supports a more separable representation space, where risk states are easier to distinguish from non-risk states. It also yields a more robust representation basis, which remains stable when the transaction network exhibits heterogeneous connectivity or when behavioral signals are partially overlapping. In this sense, deep representation learning provides more than feature compression. It offers a mechanism for extracting structured and temporally informed risk cues that are difficult to capture with rule-based or purely static modeling.

Compared with approaches that focus only on static attributes or single-step aggregation, the proposed method is more capable of capturing how risks gradually accumulate and propagate through the network. It can characterize multi-hop dependencies and latent risk chains that unfold over time, rather than treating risk as a set of isolated abnormalities. Risk identification, therefore, moves beyond post hoc explanation and becomes closer to continuous perception and dynamic assessment of risk evolution. This shift is important

for audit practice, where screening often requires ranking large populations of firms or transactions under limited resources. The results further suggest that integrating structured deep representations with dynamic risk modeling is a key direction for improving the effectiveness and practical applicability of automated audit risk screening. It can strengthen both prioritization reliability and decision consistency in complex, real-world auditing environments.

The size of the representation dimension determines the capacity and granularity of the model's latent space, directly affecting the quality of encoding enterprise transaction behavior and structural cues. Too small a dimension can lead to insufficient information compression to accommodate complex risk patterns, while too large a dimension may introduce redundancy and increase training instability. Therefore, it is necessary to observe the response pattern of the model output as the representation space capacity changes under a set of representative dimension configurations, thus providing a basis for subsequent parameter setting. The experimental results are shown in Figure 2.

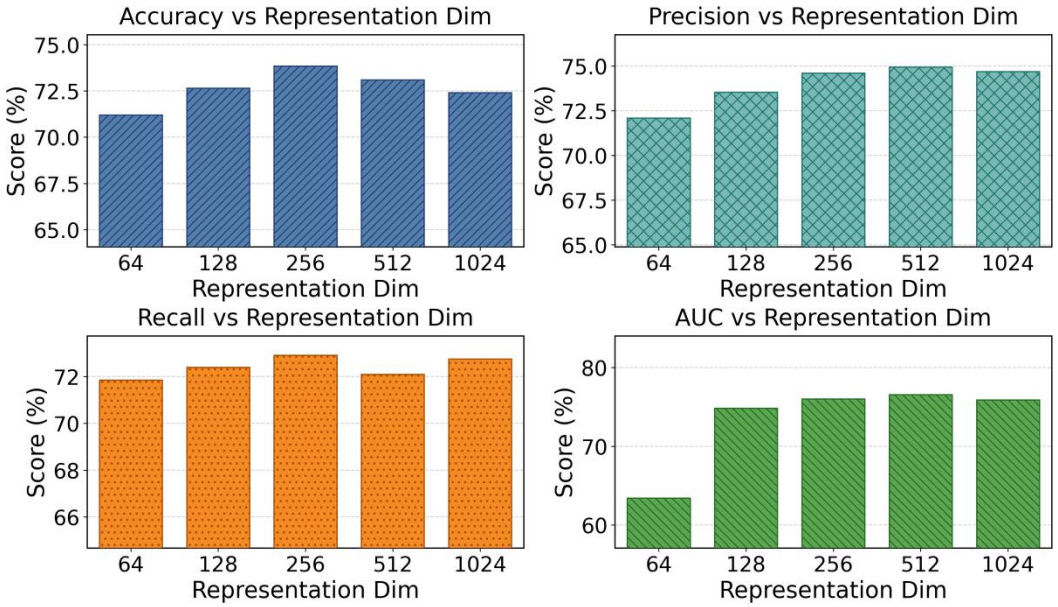


Figure 2. Indicates the effect of dimensionality on experimental results

The overall variation trend shows that representation dimensionality has a significant impact on model performance in dynamic audit risk detection tasks. This finding indicates that the capacity of the latent representation space is directly related to the model's ability to characterize complex corporate behaviors and underlying risk structures. As the representation space gradually expands, the model can absorb richer transactional semantics and structural information. Risk states, therefore, become more distinguishable in the latent space, which improves overall discriminative stability. This observation is highly consistent with the fact that audit risk is rarely triggered by a single anomaly but instead emerges from the long-term accumulation of multiple behavioral patterns.

A closer examination of the changes across different evaluation metrics reveals that their responses to representation dimensionality are not fully aligned. This suggests that dimensionality influences not only the sufficiency of information retention but also the balance among different risk discrimination priorities. Certain dimensional settings are more effective in capturing key features of risk samples, while others perform more stably in suppressing interference from redundant information. This heterogeneity reflects an inherent tradeoff in audit risk modeling, where a reasonable balance must be achieved between expressive power and representation compactness.

From the perspective of risk modeling, changes in representation dimensionality directly affect the evolution trajectory of dynamic risk states. When the representation space is insufficient to accommodate complex behavioral dependencies, risk state updates are more vulnerable to noise. This results in unstable or delayed

responses. When the representation space is overly expansive, risk signals may be diluted by irrelevant information, which weakens sensitivity to critical risk changes. A well-chosen dimensionality helps maintain the temporal coherence of risk states. It allows the risk evolution process to better reflect the actual logic of corporate operations.

Within the deep representation learning based framework for dynamic audit risk detection proposed in this study, these results confirm the foundational role of representation space design. An appropriate dimensionality not only enhances the ability to model complex audit risk patterns but also provides a stable semantic basis for subsequent dynamic state updates and risk scoring. This demonstrates that, for practical audit applications, systematic analysis and constraint of representation dimensionality are critical for ensuring model interpretability, robustness, and real-world usability.

The learning rate determines the step size and convergence rhythm of model parameter updates, directly affecting the optimization stability during representation learning and risk state modeling. To avoid excessively slow convergence or update oscillations during training, it is necessary to observe the response pattern of the model output as the optimization step size changes under a set of representative learning rate configurations. The experimental results are shown in Figure 2.

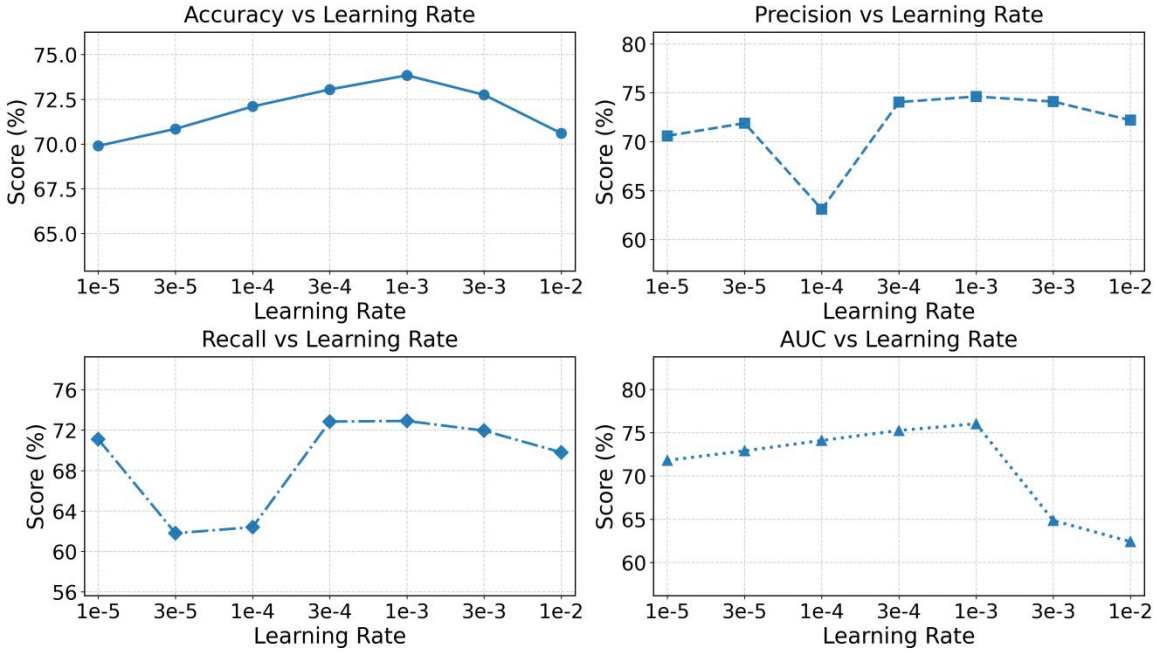


Figure 3. The impact of the learning rate on experimental results

The curve patterns indicate that the learning rate has a decisive influence on the optimization behavior of the model in dynamic enterprise audit risk detection. When the learning rate is too low, parameter updates are overly constrained. The model struggles to absorb key structural signals from transaction sequences on time. Both representation learning and risk state updates become conservative. In contrast, an excessively large learning rate tends to induce oscillatory updates. This causes unstable jumps in risk states over time and weakens the ability to continuously characterize risk evolution. Such behavior conflicts with the requirement for long-term consistency in audit risk assessment.

Across different learning rate ranges, the model outputs exhibit clear stage-dependent differences. This indicates that the optimization process does not change linearly with the learning rate. Instead, it depends on the existence of a stable convergence region. When the learning rate falls within a reasonable range, deep representations can more fully encode high-order relationships in enterprise transactions. Risk state transitions become more coherent. This leads to more reliable risk ranking and clearer decision boundaries. When the learning rate deviates from this range, the curves show pronounced fluctuations or degradation.

This reflects that the aggregation of risk signals is disrupted by noise and gradient instability, which reduces the quality of risk representations.

Differences in metric responses further suggest that the learning rate affects not only overall fitting quality but also the balance among different risk discrimination priorities. Under certain learning rate settings, the model becomes more sensitive to risk samples. It actively tracks potential risk chains. At the same time, stronger error propagation may occur, which makes decision boundaries more volatile. Other learning rate settings exhibit greater robustness. Risk state updates are smoother. However, responsiveness to short-term shocks or rare patterns may be limited. This contrast reflects a core tension in dynamic audit scenarios. Models must remain sensitive to abnormal evolution while avoiding misclassification of normal operational fluctuations as risk.

Within the proposed framework that centers on deep representation learning and risk state recursion, these observations highlight the learning rate as a key control factor linking representation quality and dynamic risk stability. An appropriate learning rate enables an optimization trajectory that balances expressiveness and stability during training. It guides the model toward latent representations that focus on generalizable structured risk patterns. It also ensures that risk states evolve in a consistent and traceable manner. Therefore, learning rate sensitivity analysis is not merely a hyperparameter tuning task. It is a necessary step for ensuring reliability and deployability in practical audit risk monitoring.

This paper also presents an experiment on the sensitivity of the proportion of missing features to accuracy, and the experimental results are shown in Figure 4.

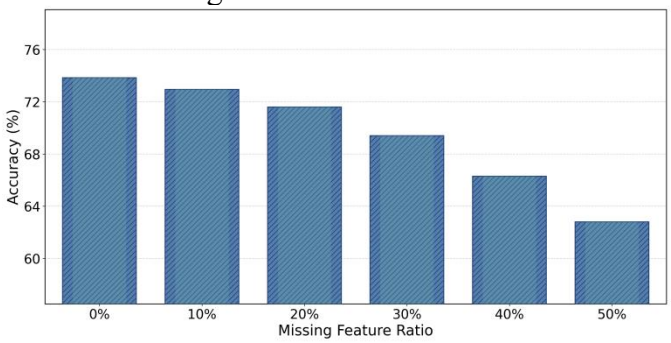


Figure 4. Sensitivity experiment of the feature missing ratio to accuracy

As the proportion of missing features increases, the stability of the model in audit risk identification is continuously challenged. This pattern indicates that deep representation learning relies to some extent on the completeness of input features. When critical information is gradually removed, the model struggles to maintain an accurate characterization of overall corporate behavior. This weakens the clear separation of risk states in the latent space. The observation is highly consistent with the uncertainty introduced by incomplete data collection in real audit settings.

A closer inspection shows that the performance degradation is not abrupt but follows a relatively smooth downward trend. This suggests that the model retains a certain level of adaptability under mild feature missingness. By jointly modeling the remaining features, the risk representations can preserve structural consistency within a limited range. Risk state updates, therefore, do not immediately collapse. As the degree of missingness increases, however, the semantic redundancy embedded in deep representations is gradually exhausted. The model's ability to abstract risk patterns then declines noticeably.

From the perspective of risk modeling mechanisms, feature missingness reduces the reliability of state propagation across time steps. When input information is insufficient to support the modeling of key transactional relationships and behavioral dependencies, risk state updates become more susceptible to noise. This manifests as delayed or shifted responses to true risk signals. Such effects impair not only risk judgments at individual time points but also accumulate over time. As a result, the dynamic risk evolution trajectory deviates from the actual logic of corporate operations.

In relation to the dynamic detection of enterprise audit risk addressed in this study, these findings indicate that the model maintains a certain degree of resilience under incomplete data conditions. However, its effectiveness depends on the availability of fundamental feature information. This highlights the importance of constraining and supplementing data quality in practical audit applications. It also suggests that combining structured deep representations with dynamic risk modeling can help mitigate the impact of partial information loss. This provides a more practically adaptive technical pathway for risk monitoring in complex audit environments.

Noise injection is used to simulate common recording errors, caliber biases, and external disturbances in enterprise audit data, which directly change the signal-to-noise ratio of the input features and the stability of representation learning. By controlling the noise intensity and observing the response pattern of the model output as the disturbance increases, the robustness boundary of the method in complex real-world environments can be evaluated. The experimental results are shown in Figure 5.

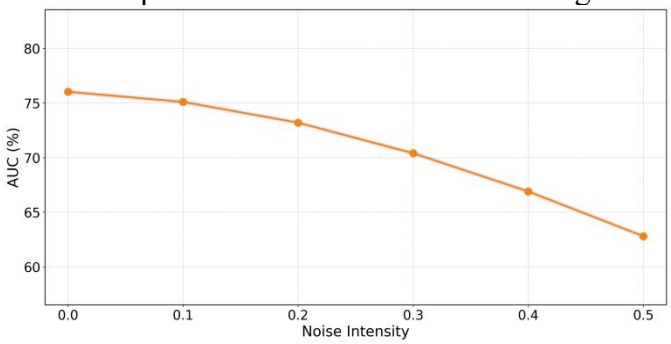


Figure 5. Noise injection intensity as a susceptibility to the AUC experiment

An examination of the overall trajectory shows that, as the intensity of noise injection increases, the model's ability to discriminate enterprise audit risk steadily declines. This indicates that deep representation learning becomes progressively disrupted when facing input perturbations. The encoding of risk-related structures and behavioral cues is weakened. Clear risk boundaries in the latent space gradually blur. This behavior closely reflects real audit settings, where noise is introduced by data collection errors or inconsistent business definitions.

From the perspective of dynamic risk modeling, noise undermines the temporal continuity of risk states. State transitions become more susceptible to random fluctuations. At low noise levels, the model can still rely on the linkage between historical states and current inputs to maintain a relatively stable risk evolution path. As perturbations intensify, the proportion of effective signals in risk state updates decreases. Risk evolution paths increasingly deviate from the true logic of corporate operations. Such deviation directly affects the reliability of ranking and identifying potential high-risk entities.

In terms of robustness, the smooth change of the curves suggests that the model does not fail immediately when noise is introduced. It exhibits a certain level of resistance to interference. This reflects that deep representation learning has some tolerance to redundancy when integrating multidimensional features. Basic risk discrimination structures can be preserved when only part of the features is disturbed. When noise accumulates to a higher level, however, redundant information is gradually eroded. The model's ability to focus on risk signals drops noticeably. The robustness boundary then becomes apparent.

With respect to the dynamic detection of enterprise audit risk examined in this study, these results indicate that noise control is critical for practical deployment. The method can maintain an overall perception of risk structure under low to moderate perturbations. Its upper performance limit remains constrained by input quality. This further highlights the necessity of combining deep representation learning with dynamic risk modeling. Structured modeling and temporal constraints can buffer part of the noise impact. They enhance model stability and usability in complex audit environments.

4. Discussion and limiting analysis

In dynamic enterprise audit risk detection tasks, model performance depends not only on structural design but also on data quality, parameter settings, and the complexity of application scenarios. The analyses in this study focus on representation dimensionality, learning rate, feature missingness, and noise perturbations. The results show that deep representation learning has clear advantages in modeling complex risk structures. Its effectiveness, however, relies on the completeness of input information and appropriate optimization configurations. When data distributions or input conditions change, the model's perception of risk states adjusts accordingly. This reflects the unavoidable uncertainty faced by audit risk modeling in real-world environments.

From the perspective of methodological applicability, the proposed dynamic detection framework demonstrates good stability under conditions of moderate complexity and controllable noise. Performance boundaries still exist under extreme feature missingness or strong perturbations. These limitations do not originate from a single module. They arise from the accumulation of errors across deep representation compression, temporal state propagation, and risk scoring. In particular, when enterprise transaction networks are highly sparse or key features remain unobservable for long periods, risk state updates may drift. This affects the accurate understanding of long-term risk evolution trends.

In addition, sensitivity to parameter selection highlights the need for caution in practical deployment. Audit scenarios differ substantially in data scale, update frequency, and risk distribution. A unified parameter configuration cannot satisfy all application requirements. Therefore, the proposed method is better viewed as a general structured risk modeling framework. In practice, it requires adaptation and constraint based on specific business characteristics. Further extension to more complex audit environments will depend on incorporating stronger data adaptive mechanisms and uncertainty modeling techniques. These advances are needed to mitigate the identified limitations and improve long-term robustness.

5. Conclusion

This study addresses the problem of dynamic identification of enterprise audit risk in complex operational environments. It proposes a unified modeling framework centered on deep representation learning. The framework maps multiple source corporate behaviors into a continuously evolving risk state space. By systematically modeling latent structural relationships and temporal dependencies in corporate activities, the method moves beyond traditional static audit analysis that relies on manual rules and isolated indicators. Audit risk is no longer treated as an isolated event. It is instead understood as a process that evolves over time and across structures. This perspective provides a more holistic and continuous analytical foundation for audit risk research.

From an application perspective, the study offers a technical pathway that aligns more closely with real business logic in audit risk management. As transaction volumes grow and business structures become increasingly complex, traditional manual audits and static analysis struggle to support continuous monitoring and early warning. By dynamically characterizing risk states, the proposed method helps auditors and regulators identify potential high-risk entities more effectively from large enterprise populations. It supports better allocation of audit resources and improves the timeliness and consistency of risk screening. This has practical significance for advancing intelligent transformation in auditing.

At a broader level, the proposed framework is not limited to enterprise audit risk scenarios. It also has potential for transfer to other complex risk management tasks. In areas such as financial supervision, compliance review, and internal control, risk similarly arises from multi-actor interactions and long-term accumulation. The emphasis on deep representations and dynamic state modeling provides a reusable analytical paradigm for these domains. It can enhance understanding of complex systemic risks and strengthen the reliability of data-driven decision-making.

Looking ahead, audit risk modeling offers substantial room for further development as enterprise data becomes richer and regulatory environments continue to evolve. On one hand, incorporating more diverse data sources and finer-grained structural modeling may further improve the characterization of concealed risk patterns. On the other hand, introducing more adaptive learning mechanisms and interpretability constraints can enhance model credibility and practical feasibility. Overall, this work establishes a solid foundation for dynamic modeling of enterprise audit risk. It also outlines a valuable direction for intelligent risk management research in related application domains.

References

[1] Motie S, Raahemi B. Financial fraud detection using graph neural networks: A systematic review[J]. Expert Systems with Applications, 2024, 240: 122156.

[2] West J, Bhattacharya M. Intelligent financial fraud detection: a comprehensive review[J]. Computers & security, 2016, 57: 47-66.

[3] Ngai E W T, Hu Y, Wong Y H, et al. The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature[J]. Decision support systems, 2011, 50(3): 559-569.

[4] Hernandez Aros L, Bustamante Molano L X, Gutierrez-Portela F, et al. Financial fraud detection through the application of machine learning techniques: a literature review[J]. Humanities and Social Sciences Communications, 2024, 11(1): 1-22.

[5] Kamuangu P. A review on financial fraud detection using ai and machine learning[J]. Journal of Economics, Finance, and Accounting Studies, 2024, 6(1): 67.

[6] Islam S, Haque M M, Karim A N M R. A rule-based machine learning model for financial fraud detection[J]. International Journal of Electrical & Computer Engineering (2088-8708), 2024, 14(1).

[7] Udayakumar R, Joshi A, Boomiga S S, et al. Deep fraud Net: A deep learning approach for cyber security and financial fraud detection and classification[J]. Journal of Internet Services and Information Security, 2023, 13(3): 138-157.

[8] Li S. Adaptive Data Mining-Based Deep Detection for Streaming Time-Series Anomaly Detection[J]., 2025, 5(6).

[9] Wang Z. A Hierarchical Encoding and Selective Attention Enhancement Framework Based on Large Language Models for Long-Document Text Classification[J], 2025, 4(11).

[10] Lai I H. Causal Aware Anomaly Detection for Modeling the Evolution of Backend Operational States in Ultra Large Scale Online Services[J], 2025, 5(7).

[11] Wang F, Verma D, Chen Y. Multi-Resource Workload Forecasting in Cloud Data Centers via Spatio-Temporal Graph Transformers[J], 2025, 5(8).

[12] Peng C C. Joint Scheduling Method for Resource Orchestration Optimization in Distributed Microservices via Hierarchical Reinforcement Learning[J], 2025, 4(9).

[13] Shu S. Adaptive Resource Scheduling in Distributed Systems via Reinforcement Learning under Dynamic Competition[J]. Artificial Intelligence and Computing Innovations, 2025, 5(7).

[14] Su J Y. Vision-Language Model-Based Emotional Understanding and Interpretable Feedback for Art Therapy Support[J], 2025, 5(7).

[15] Cui, Yiwen, et al. "FraudGNN-RL: a graph neural network with reinforcement learning for adaptive financial fraud detection." IEEE Open Journal of the Computer Society (2025).

[16] Lin, Junhong, et al. "FraudGT: a simple, effective, and efficient graph transformer for financial fraud detection." Proceedings of the 5th ACM International Conference on AI in Finance. 2024.

[17] Cheng, Zhang, et al. "Finstack-net: Hierarchical feature crossing and stacked ensemble learning for financial fraud detection." Proceedings of the 2025 2nd International Conference on Digital Economy, Blockchain and Artificial Intelligence. 2025.

[18] Abadi, Aydin, et al. "Starlit: Privacy-preserving federated learning to enhance financial fraud detection." arXiv preprint arXiv:2401.10765 (2024).
[19] Kesharwani, Abhishek, and Prashant Shukla. "FFDM− GNN: A Financial Fraud Detection Model using Graph Neural Network." 2024 International Conference on Computing, Sciences and Communications (ICCSC). IEEE, 2024.