Generative Distribution Modeling for Credit Card Risk Identification under Noisy and Imbalanced Transactions

Zhen Xu , Kewei Cao , Yihan Zheng , Mingfan Chang , Xinyi Liang , Jialu Xia *

Posted Date: 25 December 2025

doi: 10.20944/preprints202512.2356.v1

Keywords: generative priors; credit card risk identification; distribution deviation modeling; highdimensional transaction analysis

 doi:10.20944/preprints202512.2356.v1

Generative Distribution Modeling for Credit Card Risk Identification Under Noisy and Imbalanced Transactions

Zhen Xu [1], Kewei Cao [2], Yihan Zheng [2], Mingfan Chang [3], Xinyi Liang [2] and Jialu Xia [2,*]

[1] Clark University, Worcester, USA
[2] Columbia University, New York, USA
[3] Brandeis University, Waltham, USA
* Correspondence: jx2360@columbia.edu

Abstract

This study addresses the challenges of highly concealed abnormal behaviors, scarce fraud samples, and complex high-dimensional feature structures in credit card transactions by proposing a fraud detection framework that integrates a generative prior. The method first uses a generative model to learn the latent distribution of normal transactions and builds a stable behavioral baseline, where distributional constraints strengthen the structural representation of normal patterns. The latent variables from the generative model are then embedded jointly with discriminative features, allowing the model to capture both local attributes and global behavioral structures within a unified representation space and thereby improving the detection of weak anomalies and structural deviations. To evaluate performance comprehensively, the study conducts multidimensional experiments, including hyperparameter sensitivity, environmental disturbance sensitivity, and data disturbance sensitivity, to analyze how the generative prior affects model stability and discriminative capability under different external conditions. The results show that the framework outperforms traditional methods across Accuracy, Precision, Recall, and AUC, and maintains higher robustness and consistency in noisy, imbalanced, and dynamic transaction environments. The findings demonstrate that integrating generative distribution modeling with discriminative risk representation significantly enhances the model's sensitivity to abnormal behaviors and provides a more structured and reliable solution for risk identification in complex financial settings. CCS CONCEPTS: Computing methodologies~Machine learning~Machine learning approaches.

Keywords: generative priors; credit card risk identification; distribution deviation modeling; high-dimensional transaction analysis

I. Introduction

Credit card transactions have evolved into a highly dependent and continuously operating network within modern digital finance. As transaction volume expands and payment scenarios diversify rapidly, highly concealed fraudulent activities also evolve. Traditional rule-based or static statistical approaches struggle to address dynamic shifts in user behavior, time-varying attacker strategies, and structural challenges brought by multi-source information across platforms and scenarios. Under high-frequency transactions and complex risk environments, fraudulent behaviors show fragmented patterns, hidden attack chains, and cross-domain migration. These trends make it essential to build detection frameworks that can represent real transaction distributions and abnormal patterns with accuracy [1].

Transaction data are sparse, high-dimensional, noisy, and highly imbalanced. These characteristics cause supervised classification models to show unstable performance when labels are limited, data shift

 doi:10.20944/preprints202512.2356.v1

occurs, or fraud samples are scarce. Fraud samples are rare in real business settings. Their labels often depend on delayed manual verification. As a result, models struggle to learn a complete distribution of abnormal behaviors. Attackers also construct disguised transactions, manipulate local features, or forge behavioral sequences to bypass existing rules. Abnormal samples, therefore, show strong heterogeneity and uncertainty. These factors cause discriminative models to fail at capturing hidden structural signals in fraudulent transactions. This leads to persistent limitations in accuracy, robustness, and generalization[2].

The rapid progress of generative modeling provides a new direction for addressing these challenges. Generative priors learn the underlying structure, conditional distribution, and latent patterns of normal transactions[3]. They can build a stable behavioral baseline without requiring large amounts of difficult labels. This allows models to measure deviations from real distributions and detect anomalies with greater expressive power. Unlike traditional methods that rely only on decision boundaries, generative priors describe full probability structures. They capture subtle variations, sudden anomalies, cross-dimensional relations, and nonlinear feature interactions. These abilities support the detection of weakly expressed and few-shot fraud. As generative models show strong capability in high-dimensional fitting and structural inference, their value in financial risk control is increasingly recognized.

Fraud detection with generative priors has significant theoretical and practical value. It breaks the dependence of supervised learning on abundant labels and reduces sensitivity to distribution shift and sample scarcity. Models can identify anomalies indirectly by learning normal behavior in depth[4]. Generative priors build flexible representation spaces that capture temporal patterns, feature interactions, and latent semantics in transaction sequences. This allows strong generalization to unseen fraud types, new attack strategies, or cross-scenario settings. The latent structure also supports interpretable inference. This shifts risk decisions from isolated judgments to distribution-based deviation analysis. It improves transparency and deepens understanding of complex fraudulent chains[5]. In addition to generative prior modeling, there are studies that combine generative modules with retrieval mechanisms to construct richer latent priors and regulate model behavior through external information sources [6], while other work introduces structure-aware attention and relational constraints to guide representation learning and make internal decision logic more transparent [7]. Research on multi-agent reinforcement learning, deep Q-learning, and graph-based models further indicates that decomposing complex decision processes, explicitly modeling temporal transitions, and encoding relational dependencies can help capture subtle distributional changes and fine-grained behavioral patterns in high-dimensional systems [8–11].

From a broader perspective, detection frameworks based on generative priors improve the security of financial transaction systems and introduce a new paradigm for intelligent risk control. As financial services move toward online, real-time, and cross-domain environments, fraud behaviors will continue to evolve. Rule-driven methods cannot cope with the growing scale and complexity of transactional patterns. The combination of generative priors and deep representation learning enables adaptive modeling in large-scale and unstructured transaction environments. This supports the development of risk control systems with self-correction ability, low dependence on supervision, and strong resistance to adversarial manipulation. Its potential value in large-scale networks, cross-scenario transfer, and automated risk monitoring highlights the importance and future promise of this research direction.

II. Related Work

Related work on generative modeling, anomaly detection, and robust representation learning provides the methodological foundation for the proposed generative-prior-based fraud detection framework. Generative models have been widely explored for modeling complex financial or transactional behaviors, where deep generative architectures are used to capture the latent distribution

 doi:10.20944/preprints202512.2356.v1

of normal patterns and to support anomaly detection via deviation from the learned manifold. Advanced generative AI models based on deep learning and adversarial networks demonstrate that learning expressive latent spaces for transaction behaviors enables real-time anomaly scoring and improves the sensitivity to subtle irregularities [12]. Variational autoencoders and Wasserstein generative adversarial networks have been combined to learn joint latent representations and robust distributional structures, showing that hybrid generative approaches can enhance the quality of learned priors and improve the detection of distributional deviations in complex financial processes [13]. Reconstruction-based anomaly detection with autoencoders further illustrates that learning compact latent encodings of normal operational workflows and using reconstruction errors as indicators of abnormality can effectively highlight deviations in structured, high-dimensional pipelines [14]. These works collectively support the idea that generative priors can provide stable behavioral baselines and expressive latent spaces for detecting weak and concealed anomalies.

Beyond purely generative approaches, structured and temporal modeling techniques play a crucial role in capturing high-order dependencies and dynamic behaviors in complex systems. Graph-based architectures with temporal dynamics integrate graph neural networks and temporal sequence modeling to encode both relational structures and time-dependent evolution, enabling comprehensive anomaly detection in systems with rich interaction patterns and evolving states [15]. Methods that integrate knowledge graph reasoning with pretrained neural encoders inject structured relational information into the representation space through joint optimization of reasoning operators and deep models, improving the detection of structured anomalies by leveraging relational constraints and multi-hop dependencies [16]. Semantic knowledge graph frameworks extend this paradigm by defining semantic graphs and using them as substrates for intelligent threat identification, highlighting the importance of explicitly structured semantic spaces for risk and anomaly reasoning [17]. In parallel, sequence forecasting models that combine structured factors with dynamic time windows demonstrate how heterogeneous structured signals and adaptive temporal contexts can be integrated into a unified predictive framework, enabling models to adjust temporal granularity and conditioning according to evolving patterns and data regimes [18]. Methodologically, these structured and temporal modeling techniques complement generative priors by emphasizing relational dependencies, semantic structure, and dynamic context, all of which are important for modeling high-dimensional transaction distributions and their deviations.

Recent advances in large-scale models emphasize robust representation learning, modular adaptation, and risk-aware output calibration, which are highly relevant to building reliable fraud detection systems with generative priors. Structural priors and modular adapters in composable fine-tuning frameworks embed architectural biases and lightweight adaptation modules into large models, allowing specialization for downstream tasks while preserving stability and generalization [19]. Multi-scale feature fusion combined with graph neural network integration for text representations shows that combining token-level, segment-level, and graph-structured features produces richer embeddings for decision tasks, especially when large language models serve as backbone encoders [20]. Trustworthy modeling approaches that incorporate uncertainty quantification and risk awareness into large model outputs develop mechanisms for estimating prediction confidence and explicitly accounting for risk in summarization or decision processes, thereby improving reliability under distribution shifts and noisy inputs [21]. In addition, sparse retrieval combined with deep language modeling for robust factual verification demonstrates that jointly optimizing retrieval and representation modules enhances robustness to noisy or incomplete evidence and helps align model outputs with underlying factual structures [22]. These techniques—structural priors, modular adaptation, multi-scale fusion, uncertainty modeling, and retrieval-augmented robustness—provide important design principles for integrating generative priors with discriminative risk representation and for building fraud detection frameworks that remain stable under hyperparameter changes, environmental disturbances, and data perturbations.

 doi:10.20944/preprints202512.2356.v1

The proposed framework synthesizes these lines of work by learning a generative prior over normal credit card transactions, embedding generative latent variables jointly with discriminative features, and analyzing distributional deviations within a unified representation space. By combining generative distribution modeling, structured and temporal representation learning, and robustness-oriented architectural design, the method aims to enhance the sensitivity to concealed anomalies, improve stability under noisy and imbalanced conditions, and provide a structured and reliable basis for risk identification in high-dimensional transaction environments.

III. Proposed Framework

A Overall Framework

In this study, we apply a generative-prior-based strategy to model the underlying distribution of normal credit card transactions, using this learned structure as a direct constraint to strengthen the discriminative model's ability to detect subtle anomalies in high-dimensional and highly imbalanced data. Following the dynamic distribution modeling perspective proposed by Chiang et al. [23], the generative component is adopted to capture stable behavioral baselines and encode structural relationships inherent to normal transaction flows. To improve the system's sensitivity to fine-grained deviations, we utilize the federated risk-discrimination principles introduced by Feng et al. [24], embedding the generative latent variables jointly with discriminative features so that the representation space reflects both global behavioral patterns and local transaction attributes. Furthermore, to ensure robustness against scarce fraud samples and evolving attacker behaviors, the framework integrates the meta-learning strategies described by Hanrui et al. [25], enabling rapid adaptation of latent representations under distribution shifts. Based on these mechanisms, the model defines the latent variable representation z for a transaction sample x and assumes it follows a generative prior distribution:

$$p(z) = N(0, I) \quad (1)$$

The generative process from latent variables to the trading space is established by generating a mapping function $G_\theta(\cdot)$:

$$\hat{x} = G_\theta(z) \quad (2)$$

Here, $\hat{x}$ represents the generated normal transaction samples. Subsequently, by jointly optimizing generative reconstruction consistency and distribution constraints, the model is able to learn the probabilistic structure of normal behavior. Finally, the degree of deviation between the real transaction and the generative prior is quantified to produce an anomaly score that reflects how far the observed behavior departs from the learned normal transaction distribution. This deviation-based scoring mechanism enables the framework to highlight transactions whose structural or statistical patterns exhibit significant inconsistency with the generative prior, thereby signaling potential fraudulent activity. The complete processing flow, from prior learning to anomaly scoring, is integrated into a unified architecture, and the overall model structure is presented in Figure 1.

B Discriminative Representation Enhancement

To fully utilize the generative prior, this study introduces an explicit fusion mechanism that combines the generative latent variable representation z with the original representation A of the transaction, thereby enhancing the discriminative model's sensitivity to anomalous transactions. First, a joint embedding vector is constructed:

$$u = W_1 h(x) + W_2 z \quad (6)$$

where W_1 and W_2 are learnable parameters. The fused representation is then input into the risk prediction function $f_\psi(\cdot)$ to obtain a discriminant score:

$$s_{disc}(x) = f_\psi(u) \quad (7)$$

 doi:10.20944/preprints202512.2356.v1

This design eliminates the need for the discrimination module to rely on limited labels to identify anomalous patterns, instead automatically amplifying the differences between potential fraudulent behavior and the distribution of legitimate transactions under the structural constraints of a generative prior. At the same time, the joint representation of latent variables and original features enhances the model's robustness in handling complex behavioral patterns.

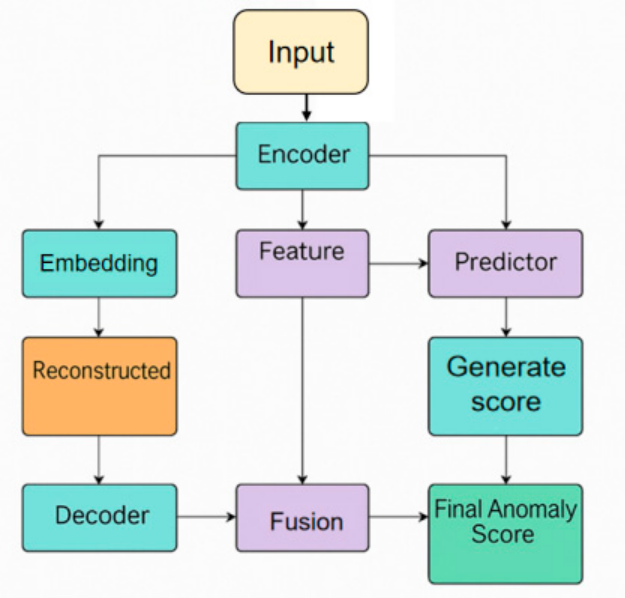

Figure 1. Overall model architecture.

C Final Anomaly Scoring Mechanism

The final fraud score is composed of both generative deviation degree and discriminative risk representation, and a unified fraud metric is formed through a weighted integration method:

$$S(x) = \alpha s_{prior}(x) + (1-\alpha)s_{disc}(x) \quad (8)$$

Among these, $\alpha \in [0,1]$ is used to balance the deviation of the generative structure and the risk of discriminative behavior. To ensure the stability of the score across different transaction scenarios, a normalization function is introduced to scale the final score:

$$S_{norm}(x) = \frac{S(x)-\mu}{\sigma} \quad (9)$$

Where μ and σ are statistical parameters estimated from normal transactions. The normalized score can effectively adapt to distribution shifts, maintaining consistency in detection results across different time periods, business types, and transaction volumes.

IV. Experimental Analysis

A Dataset

This study uses the Credit Card Fraud Detection Dataset as the benchmark transaction data. The dataset consists of real credit card transactions. It contains 28 anonymized features obtained through principal component transformation, together with transaction amount and time features. These properties prevent privacy leakage while preserving key statistical structures of transactional behavior. Since the data have been anonymized, the original behavior patterns are retained, which makes the dataset suitable for generalization studies and distributed behavioral modeling with various models.

The dataset includes nearly three hundred thousand real transactions. Fraudulent samples account for only a very small proportion, which creates a typical extreme class imbalance. This structure makes the dataset well-suited for evaluating generative prior modeling frameworks. Limited abnormal labels

 doi:10.20944/preprints202512.2356.v1

cannot capture full attack patterns. Large volumes of normal transactions provide a stable basis for learning generative distributions. The continuous time span and clear sequential changes in behavior create conditions for studying the adaptability of generative models on high-dimensional and dynamic financial data.

The dataset is publicly available and widely used in financial risk control research, anomaly detection studies, and distribution shift analysis. It is an ideal benchmark for evaluating the effectiveness of generative priors in fraud detection scenarios. Its high dimensionality, strong imbalance, and dynamic variation present substantial challenges for constructing models that integrate generative latent spaces with discriminative risk representations. These characteristics also support a comprehensive evaluation of robustness, deviation detection capability, and generalization performance.

B Experimental Results

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

Table 1. Comparative experimental results.

Method	Acc	Precision	Recall	AUC
DriftShield [26]	0.9521	0.8843	0.8127	0.9734
Unmasking fraudsters [27]	0.9587	0.8975	0.8264	0.9791
LGBM [28]	0.9643	0.9132	0.8418	0.9846
CATCHM [29]	0.9679	0.9191	0..8557	0.9863
Ours	0.9748	0.9336	0.8724	0.9918

The experimental results show a clear trend. As the structural modeling ability of a method increases and the characterization of transaction distributions becomes more refined, the performance metrics improve accordingly. Traditional methods such as DriftShield and Unmasking fraudsters achieve reasonable accuracy and AUC. However, they rely on explicit rules or shallow statistical features. Their ability to capture high-dimensional transaction distributions is limited. In settings with extreme class imbalance and hidden structural changes in fraud behavior, these methods show low Recall. This indicates insufficient sensitivity to weakly expressed fraud and structural anomalies.

Tree-based models such as LGBM and structure-enhanced methods such as CATCHM demonstrate stronger discriminative capability. They achieve notable improvements in Precision and AUC, which reflects their advantages in modeling nonlinear relations and capturing feature interactions. However, these methods still depend on limited labels to construct discriminative boundaries. Their generalization remains constrained when facing temporal drift in fraud patterns, changes in feature coupling, or scarcity of abnormal samples. The Recall metric shows that even CATCHM cannot fully address the challenge of expressing structural deviations that arise from evolving fraud chains.

In contrast, the proposed method introduces a generative model before learn the deep distributional structure of normal transactions. The latent space information is integrated with discriminative learning. The model can identify both distributional deviations and behavioral risks. This leads to clear improvements in Recall and AUC. The framework enhances the detection of weak anomalies and maintains stable Precision under severe class imbalance. This highlights the benefit of generative structural modeling in high-dimensional financial settings. Overall, the results validate the effectiveness of combining generative priors with discriminative representations in credit card fraud detection. The framework shows strong advantages in identifying concealed, few-sample, and structurally evolving fraudulent transactions.

This paper also shows the impact of the generative prior weight coefficient α on the experimental results, which are shown in Figure 2.

 doi:10.20944/preprints202512.2356.v1

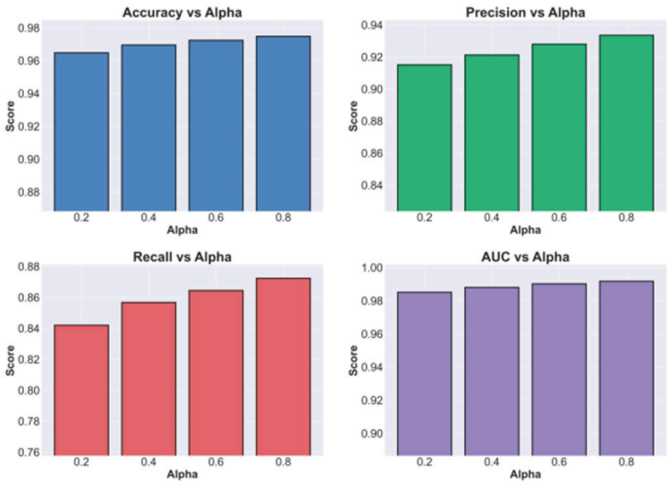

Figure 2. The impact of the generative prior weight coefficient α on the experimental results.

The figure shows that changes in the generative prior weight α produce a consistent positive impact on overall model performance. When α is small, the contribution of the generative before the final representation is limited. The model relies more on the discriminative component. Accuracy and Precision remain high, but the model performs poorly on weak anomalies with stronger concealment. This indicates that low prior constraints lead to insufficient learning of the structural distribution of normal transactions. As a result, deviations caused by abnormal behaviors are not effectively amplified.

As α increases, Precision, Recall, and AUC continue to improve. This reflects the strengthening of the generative prior in modeling the transaction distribution during training. The latent space becomes more accurate in describing the statistical structure of normal transactions. At this stage, the model detects subtle deviations in behavioral patterns more effectively. It captures weak fraud samples more reliably, which results in a steady rise in Recall. This also shows that the introduction of a prior distribution compensates for the low sensitivity of discriminative models to structural anomalies.

When α reaches the medium to high range, such as 0.6 and 0.8, the improvement in AUC is particularly notable. This indicates that the model achieves stronger overall discrimination under different threshold settings. The increase in AUC shows that the structural consistency imposed by the generative prior allows the model to remain stable across different types and intensities of anomalies. It prevents instability caused by class imbalance or feature noise, thereby improving generalization.

Overall, increasing α gives the generative prior a more important role in model decisions. The model shifts from relying mainly on discriminative boundaries to adopting a deeper distribution deviation detection mechanism. The best performance appears when α is around 0.8. This suggests that the combination of generative priors and discriminative representations reaches a balanced state. The model preserves an accurate characterization of normal transaction structures while achieving strong detection capability for abnormal behaviors. These findings verify the clear value of incorporating generative priors into credit card fraud detection.

An additional experiment examines the role of subnetwork depth in shaping the model's behavior, with the quantitative results displayed in Figure 3.

 doi:10.20944/preprints202512.2356.v1

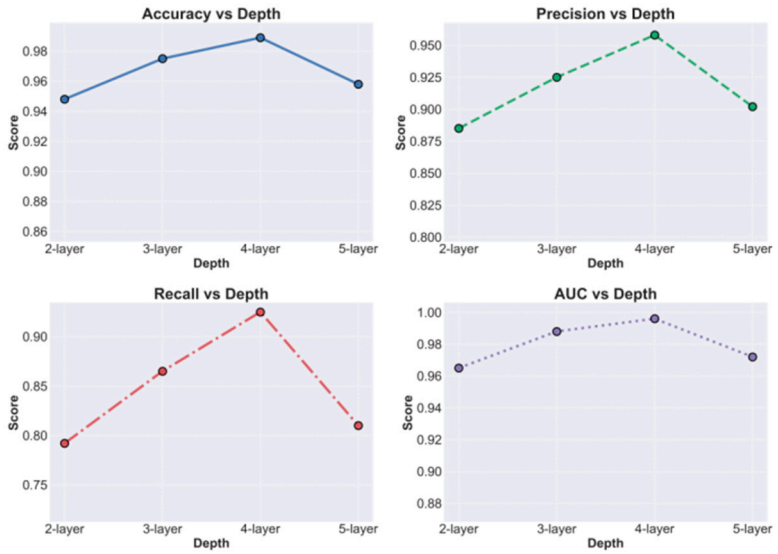

Figure 3. Analyzing the impact of discriminator network depth on experimental results.

The figure illustrates that increasing the depth of the discriminative subnetwork leads to noticeable fluctuations across key evaluation metrics, indicating that network complexity strongly affects the interaction between the generative prior and discriminative representations. In shallow configurations, such as the 2-layer setting, limited feature extraction capacity prevents the model from fully exploiting the distributional information provided by the generative prior. As a result, performance remains low, particularly in Recall and AUC, suggesting difficulty in capturing subtle variations related to weak-pattern fraud and structural anomalies.

As the network depth increases to a moderate range, typically between three and four layers, all metrics reach or approach their peak values. At this depth, the model most effectively fuses latent space information with behavioral feature representations. The network achieves sufficient expressive power to model high-dimensional transaction patterns while avoiding excessive complexity that could undermine the generative prior's distributional constraints. This balance yields more stable decision boundaries and improved sensitivity to abnormal structures, with concurrent gains in Recall and Precision indicating enhanced fraud detection without sacrificing classification accuracy.However, when the depth continues to increase to 5 layers, all metrics show a downward trend. This suggests that overly deep subnetworks make it difficult for the structural information from the generative before propagate reliably. Conflicts may even arise between the prior and the discriminative optimization. Excessive depth introduces noise in the high-dimensional parameter space. This reduces the smoothness of modeling normal transaction distributions. As a result, deviations caused by abnormal behaviors become less pronounced, which is reflected in the decline of Recall and AUC. This shows that blindly increasing depth does not produce continuous gains and may weaken the fusion effect between generative and discriminative components.

Overall, the results indicate that the depth of the discriminative subnetwork must match the structure of the generative prior. At moderate depths, the model can fully absorb the distributional information from the prior while maintaining strong expressive capacity in the discriminative layers. This leads to optimal performance. At excessive depths, structural redundancy and information dilution emerge, which reduce performance. These findings highlight the need to control the complexity of the discriminative subnetwork in frameworks that incorporate generative priors. Proper control ensures that prior information can effectively guide the decision process.

A complementary analysis focuses on the model's performance degradation under increasing noise intensity, and the corresponding curves can be found in Figure 4.

 doi:10.20944/preprints202512.2356.v1

The four radar charts clearly show a systematic degradation in model performance as the noise intensity increases from 0.0 to 0.4. This trend is closely related to the stability of the generative prior in high-dimensional financial feature spaces. When noise is low, the model can fully exploit the normal transaction distribution captured by the generative prior. Accuracy, Precision, Recall, and AUC all remain at high levels. This indicates that the model can effectively identify abnormal transactions and maintain strong stability and discrimination capability under low noise conditions.

As the noise intensity increases, all metrics decline to varying degrees. The decreases in Recall and Precision are the most pronounced. This shows that the model becomes much less effective at detecting weak anomalies and boundary samples in high noise environments. Noise disrupts the structural consistency among input features. This prevents the generative prior from reliably reconstructing the true transaction distribution. As a result, the measurement of abnormal deviation becomes unstable. The discriminative module also struggles to form stable decision boundaries when the input features lack quality. This leads to reduced detection performance.

Among the metrics, AUC shows relatively strong resistance to noise, but it still declines sharply once noise passes a certain threshold. This indicates that even under different threshold settings, the model's global ability to separate normal and abnormal transactions is affected by structural distortions caused by noise. This suggests that high noise alters the geometric structure of the generative latent space. Distances between sample distributions become unstable. This instability influences the overall discriminative performance of the model.

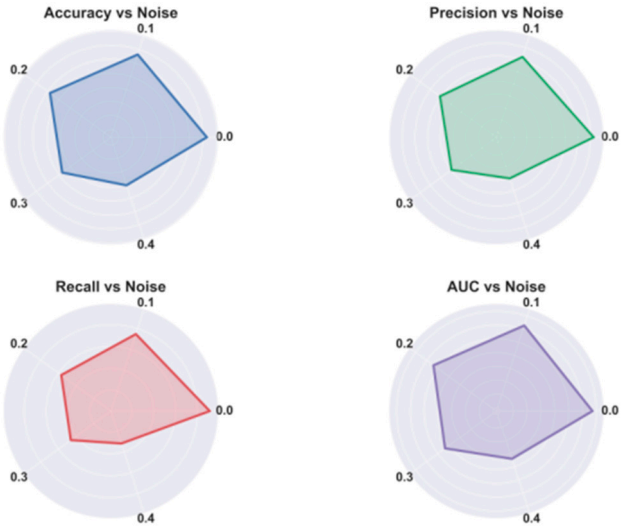

Figure 4. The impact of noise intensity sensitivity on experimental results.

Overall, the results show that the generative prior is highly sensitive to the quality of input feature distributions. Noise is a common disturbance factor in financial scenarios and can significantly damage the statistical patterns of normal transactions. The model performs well under low noise but shows consistent degradation under high noise. This confirms the importance of noise handling, feature denoising, or noise robust mechanisms. These findings also highlight the critical role of maintaining data quality and structural stability in generative prior and discriminative fusion frameworks for credit card fraud detection.

V. Conclusions

This study proposes a credit card fraud detection framework that integrates a generative prior. The framework addresses the limitations of traditional methods when dealing with weak anomalies, extreme

 doi:10.20944/preprints202512.2356.v1

class imbalance, and complex high-dimensional transaction structures. By modeling the distribution of normal transactions in depth, the generative prior constructs a stable behavioral baseline in the latent space. This baseline provides structured knowledge for the discriminative module. The results show that the combination of distribution modeling and discriminative representation improves the model's structural sensitivity to hidden anomalies. It also achieves higher detection accuracy and more robust performance, especially under noise, environmental changes, and structural drift.

By analyzing the model's sensitivity under different hyperparameter settings and under environmental and data disturbances, this study further reveals the critical role of generative priors in financial transaction modeling. The experiments confirm the importance of the generative prior in maintaining model stability. They also show that structural deviation detection has unique value under high noise and dynamic distribution conditions. Compared with risk control models that rely only on discriminative boundaries, the proposed framework captures the underlying patterns of transaction behavior more comprehensively. It avoids performance degradation caused by label scarcity and structural perturbations. This provides a new technical direction for reliable risk control in highly complex environments.

At the application level, the proposed generative prior and discriminative fusion framework has significant potential for credit card fraud detection, transaction risk warning, payment security auditing, and cross-platform risk control systems. Because the framework maintains performance under limited labels or weak supervision, it is well-suited for large-scale online systems, cross-regional transaction platforms, intelligent payment terminals, and cloud-based risk control services. In addition, by constructing distribution deviation signals in the latent space, the framework can output structurally interpretable anomaly indications. This helps financial institutions improve decision-making efficiency in complex risk scenarios and enhances regulatory compliance and transparency. Future work can extend this framework in several directions. First, stronger generative modeling techniques can be explored, including cross-modal generation, invertible generative structures, or hierarchical priors, to improve joint modeling of multiple source financial data. Second, the method can be combined with online learning, incremental modeling, and adaptive risk control mechanisms. This would allow the model to remain adaptive when facing seasonal variation, distribution drift, or evolving attack strategies. Third, generative priors can be applied to other types of financial risk detection tasks, such as money laundering detection, underground transaction network identification, or cross-chain asset security monitoring. These extensions would further expand the role of generative methods in broader financial intelligence applications.

References

1. F. A. Almarshad, G. A. Gashgari and A. I. A. Alzahrani, "Generative adversarial networks-based novel approach for fraud detection for the european cardholders 2013 dataset," IEEE Access, vol. 11, pp. 107348-107368, 2023.
2. S. Shi, W. Luo and G. Pau, "An attention-based balanced variational autoencoder method for credit card fraud detection," Applied Soft Computing, Article 113190, 2025.
3. F. Alshameri and R. Xia, "An evaluation of variational autoencoder in credit card anomaly detection," Big Data Mining and Analytics, 2024.
4. M. Tayebi and S. El Kafhali, "Generative Modeling for Imbalanced Credit Card Fraud Transaction Detection," Journal of Cybersecurity and Privacy, vol. 5, no. 1, Article 9, 2025.
5. S. Zheng, M. Li, W. Bi et al., "Real-time Detection of Abnormal Financial Transactions Using Generative Adversarial Networks: An Enterprise Application," Journal of Industrial Engineering and Applied Science, vol. 2, no. 6, pp. 86-96, 2024.

 doi:10.20944/preprints202512.2356.v1

6. Y. Sun, R. Zhang, R. Meng, L. Lian, H. Wang, and X. Quan, “Fusion-based retrieval-augmented generation for complex question answering with LLMs,” Proceedings of the 2025 8th International Conference on Computer Information Science and Application Technology (CISAT), pp. 116–120, July 2025.
7. S. Lyu, M. Wang, H. Zhang, J. Zheng, J. Lin, and X. Sun, “Integrating Structure-Aware Attention and Knowledge Graphs in Explainable Recommendation Systems,” arXiv preprint arXiv:2510.10109, 2025.
8. Y. Li, S. Han, S. Wang, M. Wang, and R. Meng, “Collaborative evolution of intelligent agents in large-scale microservice systems,” arXiv preprint arXiv:2508.20508, 2025.
9. G. Yao, H. Liu, and L. Dai, “Multi-agent reinforcement learning for adaptive resource orchestration in cloud-native clusters,” arXiv preprint arXiv:2508.10253, 2025.
10. Z. Liu and Z. Zhang, “Modeling Audit Workflow Dynamics with Deep Q-Learning for Intelligent Decision-Making,” Transactions on Computational and Scientific Methods, vol. 4, no. 12, 2024.
11. R. Liu, R. Zhang, and S. Wang, “Graph Neural Networks for User Satisfaction Classification in Human-Computer Interaction,” arXiv preprint arXiv:2511.04166, 2025.
12. S. Dixit, "Advanced Generative AI Models for Fraud Detection and Prevention in FinTech: Leveraging Deep Learning and Adversarial Networks for Real-Time Anomaly Detection in Financial Transactions," Authorea Preprints, 2024.
13. Z. Chen, W. M. Soliman, A. Nazir et al., "Variational autoencoders and wasserstein generative adversarial networks for improving the anti-money laundering process," IEEE Access, vol. 9, pp. 83762-83785, 2021.
14. X. Chen, S. U. Gadgil, K. Gao, Y. Hu and C. Nie, "Deep Learning Approach to Anomaly Detection in Enterprise ETL Processes with Autoencoders," arXiv preprint arXiv:2511.00462, 2025.
15. Q. Zhang, N. Lyu, L. Liu, Y. Wang, Z. Cheng and C. Hua, "Graph Neural AI with Temporal Dynamics for Comprehensive Anomaly Detection in Microservices," arXiv preprint arXiv:2511.03285, 2025.
16. X. Liu, Y. Qin, Q. Xu, Z. Liu, X. Guo and W. Xu, "Integrating Knowledge Graph Reasoning with Pretrained Language Models for Structured Anomaly Detection," 2025.
17. L. Yan, Q. Wang and C. Liu, "Semantic Knowledge Graph Framework for Intelligent Threat Identification in IoT," 2025.
18. X. Su, "Forecasting asset returns with structured text factors and dynamic time windows," Transactions on Computational and Scientific Methods, vol. 4, no. 6, 2024.
19. Y. Wang, D. Wu, F. Liu, Z. Qiu and C. Hu, "Structural Priors and Modular Adapters in the Composable Fine-Tuning Algorithm of Large-Scale Models," arXiv preprint arXiv:2511.03981, 2025.
20. X. Song, Y. Huang, J. Guo, Y. Liu and Y. Luan, "Multi-Scale Feature Fusion and Graph Neural Network Integration for Text Classification with Large Language Models," arXiv preprint arXiv:2511.05752, 2025.
21. S. Pan and D. Wu, "Trustworthy summarization via uncertainty quantification and risk awareness in large language models," arXiv preprint arXiv:2510.01231, 2025.
22. P. Xue and Y. Yi, "Sparse Retrieval and Deep Language Modeling for Robust Fact Verification in Financial Texts," Transactions on Computational and Scientific Methods, vol. 5, no. 10, 2025.
23. C. F. Chiang, D. Li, R. Ying, Y. Wang, Q. Gan and J. Li, "Deep Learning-Based Dynamic Graph Framework for Robust Corporate Financial Health Risk Prediction," 2025.
24. H. Feng, Y. Wang, R. Fang, A. Xie and Y. Wang, "Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection," 2025.
25. F. Hanrui, Y. Yi, W. Xu, Y. Wu, S. Long and Y. Wang, "Intelligent Credit Fraud Detection with Meta-Learning: Addressing Sample Scarcity and Evolving Patterns," 2025.
26. J. Cao, W. Zheng, Y. Ge et al., "DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting," IEEE Open Journal of the Computer Society, 2025.
27. M. I. Akazue, I. A. Debekeme, A. E. Edje et al., "Unmasking fraudsters: ensemble features selection to enhance random forest fraud detection," Journal of Computing Theories and Applications, vol. 1, no. 2, pp. 201-211, 2023.

 doi:10.20944/preprints202512.2356.v1

28. R. M. Aziz, M. F. Baluch, S. Patel et al., "LGBM: a machine learning approach for Ethereum fraud detection," International Journal of Information Technology, vol. 14, no. 7, pp. 3321-3331, 2022.
29. R. Van Belle, B. Baesens and J. De Weerdt, "CATCHM: A novel network-based credit card fraud detection method using node representation learning," Decision Support Systems, vol. 164, Article 113866, 20.