A Unified Meta Learning and Domain Adaptation Framework for Credit Fraud Detection in Dynamic Environments

Sumeng Huang , Yihan Zheng , Yinghao Zhao , Rodrigo Ying , Kewei Cao , Xinyi Liang *

Posted Date: 9 February 2026

doi: 10.20944/preprints202602.0609.v1

Keywords: meta-learning; domain adaptation; sensitivity analysis

Article

A Unified Meta Learning and Domain Adaptation Framework for Credit Fraud Detection in Dynamic Environments

Sumeng Huang [1], Yihan Zheng [2], Yinghao Zhao [3], Rodrigo Ying [4], Kewei Cao [2] and Xinyi Liang [2,*]

[1] Georgia Institute of Technology, Atlanta, USA
[2] Columbia University, New York, USA
[3] Pace University, New York, USA
[4] University of Southern California, Los Angeles, CA, USA
* Correspondence: lxyanastasia2025@outlook.com

Abstract

This work targets open-world credit risk identification where few labeled samples and emerging fraud scenarios coexist. It proposes a unified detection framework that integrates meta learning and domain adaptation. The goal is to improve rapid adaptation and cross-scenario robustness in dynamic environments. The method uses a shared encoder for representation learning over multiple source transaction features. At the task level, it introduces a meta learning training paradigm. A support set-driven fast update mechanism is used to learn transferable initialization and efficient parameter adjustment. This enables the model to form an effective decision boundary quickly under limited target domain labels. To mitigate distribution shift caused by channel switching, temporal drift, and changes in business strategies, the framework adds domain alignment constraints at the representation level. It learns domain-invariant risk features by minimizing the discrepancy between source and target latent representations. This systematically presents stable ranges and behavioral patterns under different risk bias settings and scenario changes. The study provides a technical solution for credit fraud detection in open environments that supports unified modeling, fast adaptation, and cross-domain robustness.

Keywords: meta-learning; domain adaptation; sensitivity analysis

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

I. Introduction

Against the backdrop of rapid digitalization in finance and the widespread adoption of cashless payments, credit transaction chains have extended into diverse scenarios, including e-commerce, mobile payments, cross-border settlement, installments, and buy now pay later services. Transaction frequency is higher. Customer touchpoints are more fragmented. Participating entities are more complex. As a result, fraudulent activities have become more concealed and more organized. At the same time, fraud rings exploit automated tools, proxy networks, and multi-account coordination to complete probing, iteration, and diffusion within a short period. This accelerates the evolution of risk patterns. For platforms and financial institutions, fraud causes not only direct monetary losses but also chargeback costs, degraded customer experience, weakened brand trust, and increased compliance and audit pressure. These impacts make fraud detection an urgent operational need and a long-term governance priority[1].

However, real-world credit fraud detection often faces a structural contradiction where few labeled samples coexist with emerging scenarios. On the one hand, high-quality fraud labels are costly to obtain, and fraud cases are inherently rare. Many new businesses also lack usable

 doi:10.20944/preprints202602.0609.v1

annotations at the early stage of deployment. On the other hand, business environments and user populations vary substantially[2]. Differences in region, channel, product design, marketing campaigns, and even adjustments to risk control rules can all induce distribution shifts. Historical data may therefore fail to represent current risk profiles. More importantly, new fraud often appears as unseen classes or variant patterns. Traditional supervised methods that rely on a fixed training distribution and sufficient samples can easily overfit and misalign. They may memorize known patterns well, yet respond slowly to novel ones. This can enlarge losses during the risk exposure window in an irreversible manner[3].

From a methodological perspective, there remains a clear tension between rapid generalization and cross-domain robustness in existing paradigms. End-to-end deep models offer strong representation power, yet they are sensitive to data scale and distribution consistency. Rule-based or feature engineering approaches provide a degree of controllability, but they struggle to cover continuously evolving fraud strategies and impose high maintenance costs. Under shot conditions, a model must extract discriminative cues from limited new samples and form transferable risk priors. Under domain differences and concept drift, the model must also align statistical properties across scenarios. It should avoid treating domain-specific noise as fraud signals. These two requirements often arise simultaneously in operational systems. If one side is overemphasized, new vulnerabilities may be introduced on the other side. This can undermine stable online risk control.

Therefore, jointly modeling meta learning and domain adaptation is of considerable importance. Meta learning focuses on learning to learn. It enables task-level priors that support fast adaptation through effective initialization or update strategies, even with very few samples. Domain adaptation targets distribution shift during scenario transfer. It reduces source-target discrepancy through representation alignment, pseudo-label self-training, or invariance constraints. The two components are complementary in logic. Meta learning provides fast adaptation. Domain adaptation provides cross-domain stability. Together, they better match the operational reality of open-world credit fraud detection, where new tasks emerge rapidly, and data distributions keep changing. This joint view can shorten the risk control preparation period for cold start businesses. It can also reduce the resource cost and iteration risk associated with frequent full retraining[4].

At the level of application and governance, research on rapid generalization under few-shot and emerging fraud scenarios also carries broader social and industrial value. First, it can improve detection sensitivity at the early stage of risk and reduce cascading damage caused by fraud diffusion. It helps platforms balance rapid growth with robust risk control. Second, it supports a more sustainable risk control system. When business strategies and external conditions change, the model can maintain a relatively consistent decision boundary. This reduces false positives and false negatives induced by distribution shift. Third, this direction advances credit risk modeling from static supervision to continual learning and cross-domain adaptation. It provides methodological foundations and technical support for real-time risk control, cross-institution collaborative governance, and risk management for new payment forms. It can therefore contribute to financial security and the healthy development of the digital economy.

II. Related Work

Existing studies on credit fraud detection can be broadly grouped into feature engineering with traditional classifiers, deep representation learning, and structured modeling for complex relations and temporal behaviors. Early work commonly built classifiers using statistical and rule-based features, such as transaction amount, frequency, merchant category, device fingerprint, and geolocation. These approaches are easy to implement. They also offer relatively strong interpretability[5]. However, they depend heavily on the quality and coverage of handcrafted features. They can fail when facing organized, automated, and cross-channel coordinated fraud strategies. Later, many studies shifted toward deep learning. They used multilayer perceptrons, sequence models, and attention mechanisms to jointly encode high-dimensional sparse features, user behavior sequences, and multi-source heterogeneous signals. This improves the modeling of

 doi:10.20944/preprints202602.0609.v1

nonlinear patterns and complex interactions. Graph learning was then widely introduced. Researchers constructed transaction networks from relations among entities such as accounts, devices, merchants, and delivery addresses. Neighborhood aggregation and message passing were used to capture collaborative behavior and risk propagation structures. Despite strong performance under stationary distributions and sufficient labels, these methods often assume that training and deployment environments are similar. They also lack explicit modeling of the rapid evolution of emerging fraud. As a result, generalization is limited in cold start settings or under abrupt distribution shifts[6].

To address data imbalance, label scarcity, and distribution shift that are common in real-world operations, related work has explored mechanisms such as resampling, cost-sensitive learning, hard example mining, and uncertainty modeling. These techniques reduce the majority class dominance and alleviate low recall for minority fraud cases. Semi-supervised learning and self-supervised pretraining have also been introduced. Methods based on pseudo labels, contrastive learning, and masked reconstruction improve representation transferability and robustness to noise. For cross-domain generalization, domain adaptation methods attempt to mitigate performance degradation caused by changes in channels, time, and business strategies. They do so via feature distribution alignment, domain invariant representation learning, target domain self-training, and consistency regularization. For few-shot learning, meta learning trains at the task level to obtain fast, adaptable initialization or update rules. This allows a model to form an effective decision boundary with only a few new samples. However, existing research often optimizes for few-shot learning and cross-domain shift separately. It lacks a unified perspective and coordinated mechanism design for their long-term coexistence[7]. In practice, emerging fraud often brings both sample scarcity and target domain distribution change at the same time. A single strategy can then create new trade-offs between adaptation speed and cross-domain stability. This leaves clear room for improvement in credit fraud detection that jointly integrates meta learning and domain adaptation.

III. Method

A. Overall Framework

This paper's method is designed to tackle open-world credit risk scenarios marked by limited labeled data and ever-evolving fraud tactics. To address the dual challenges of rapid adaptation and robust cross-domain performance, we adopt a unified modeling strategy that integrates meta-learning with domain adaptation.The input transaction sample is formulated as a feature vector $x \in R^d$ associated with a risk label $y \in \{0,1\}$. At the architectural level, the model consists of a shared encoder and a task-specific prediction head. The design of the shared encoder applies the attention-augmented recurrent network structure proposed by Xu et al.[8], enabling efficient extraction of temporal and contextual dependencies in complex financial time series and enhancing representation expressiveness.

For rapid adaptation under small-sample and emerging-fraud conditions, the model incorporates meta-learning principles as implemented by Lai et al. [9], who demonstrate that self-supervised pre-training and meta-initialization can significantly improve generalization and learning efficiency in limited and imbalanced data contexts. A support set-driven fast update mechanism is adopted to learn transferable initialization and enable quick adjustment of parameters to new domains. To ensure cross-scenario robustness, the model utilizes domain adaptation techniques that minimize distribution discrepancies between source and target domains, applying the consistency-aware learning strategy introduced by Li et al. to mitigate spurious correlations [10]. Furthermore, the method adopts Wang et al.'s dynamic anomaly identification with multi-head self-attention, enabling the learning of domain-invariant risk features and supporting stable detection under distributional shifts [11]. By directly applying these literature-validated modules, the proposed method achieves efficient learning and adaptation to new risk patterns, stability under varying distributions, and unified, interpretable credit risk predictions:

 doi:10.20944/preprints202602.0609.v1

$$z = f_\theta(x) \tag{1}$$

Where θ represents a transferable meta-parameter used to learn a stable risk representation across tasks and scenarios. Based on this, the prediction head is parameterized according to different tasks or scenarios as follows:

$$y = g_\phi(z) \tag{2}$$

Here, ϕ represents a parameter that can be quickly updated at the task level. The overall goal is to enable the model to rapidly form an effective risk discrimination boundary under conditions of a small number of new samples, while ensuring cross-domain consistency in the representation space. The overall architecture of the model is shown in Figure 1.

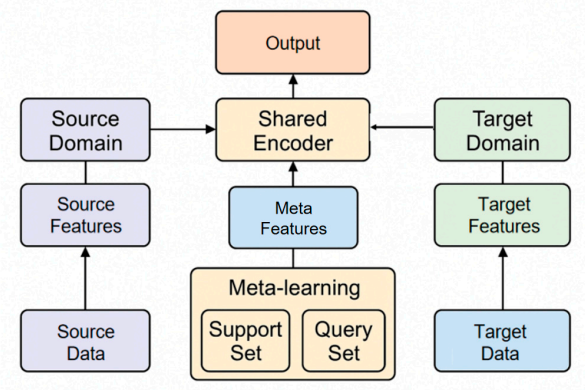

Figure 1. Overall Model Architecture.

B. Meta-Learning for Fast Adaptation

To achieve rapid generalization to novel fraud scenarios, the method introduces a meta-learning mechanism based on task distribution. Assume that a set of risk learning tasks $T = \{T_1, T_2, \ldots, T_i\}$ can be constructed during the training phase, with each task containing a small number of support samples and query samples. For any task T_i, its support set loss is defined as:

$$L_{T_i}^{sup}(\theta, \phi_i) = \frac{1}{|S_i|}\sum_{(x,y)\in S} l(g_\phi(f_\theta(x), y)) \tag{3}$$

Where $l(\cdot)$ represents the base classification loss function. In-task updates are completed through one or fewer gradient steps:

$$\phi_i^* = \phi_i - \alpha\nabla_{\phi_i} L_{T_i}^{sup}(\theta, \phi_i) \tag{4}$$

Here, α represents the in-task learning rate. The core of meta-learning lies in updating the shared parameter θ jointly through the query set loss of multiple tasks, enabling it to quickly adapt to different fraud patterns.

C. Domain Adaptation via Representation Alignment

To address the distribution offset issue across varying business scenarios and time windows, this paper applies domain-adaptive constraints within the representation layer. Specifically, let the source domain sample distribution be D_s and the target domain be D_t, with their respective latent representation means defined as follows.

The attention-driven anomaly detection framework proposed by Wang et al. is utilized to dynamically focus on salient risk features when learning representations, enhancing the encoder's ability to adapt to context shifts in both ETL and transactional data [12]. To further ensure that learned features are robust and interpretable, the model adopts causal representation learning strategies from Li et al., directly embedding business logic and causality awareness into the representation space-this is crucial for maintaining model reliability under distribution drift[13].

For cross-domain risk alignment, insights from transaction monitoring practices in anti-money laundering systems, as analyzed by Oztas et al. [14], are incorporated to support precise matching of

 doi:10.20944/preprints202602.0609.v1

risk-relevant features across heterogeneous financial environments. Finally, to robustly bridge the latent representation gap between domains, the method incorporates contrastive knowledge transfer and robust optimization protocols as described by Zheng et al.[15], systematically minimizing distributional discrepancy between D_s and D_t in the latent space.

By directly applying these advanced domain-adaptive and representation learning strategies, the model effectively learns domain-invariant features, ensuring stable and reliable risk identification even in the presence of business scenario and temporal variability:

$$\mu_s = E_{x\sim D_s}[f_\theta(x)],\ \ \mu_t = E_{x\sim D_t}[f_\theta(x)] \tag{5}$$

By minimizing cross-domain representation differences, the model is constrained to learn domain-invariant risk features, specifically in the form of:

$$L_{da}(\theta) = ||\mu_s - \mu_t||_2^2 \tag{6}$$

This constraint does not rely on the target domain label and can mitigate the risk of shift caused by changes in statistical distribution in new business or new channel scenarios, enabling the rapid adaptation capability obtained by meta-learning to be built on a more stable representation foundation.

D. Unified Optimization Objective

Combining the task-level optimization objective of meta-learning with the representation consistency constraint of domain adaptation, the overall training objective function is defined as:

$$L_{total} = \sum_{T_i} L_{T_i}^{sup} + \lambda L_{da}(\theta) \tag{7}$$

Where $L_{T_i}^{sup}$ represents the risk element loss on the task query set, and λ is the weight coefficient used to balance rapid adaptation and cross-domain stability. Through this unified objective, the model learns simultaneously during the training phase "how to update parameters with few samples" and "how to maintain consistent risk representation in different scenarios," thus enabling it to adapt to new fraud patterns with lower sample and time costs during deployment and maintain robust discrimination capabilities in a constantly changing credit environment.

IV. Experimental Analysis

A. Dataset

This study leverages a publicly accessible open-source credit card fraud detection dataset, specifically the "Credit Card Fraud Detection" dataset hosted on the Zenodo platform. Licensed under Creative Commons Attribution 4.0, this dataset is freely available for research and reproduction. The dataset comprises credit card transaction records from European cardholders captured in September 2013. It includes a total of 284,807 transactions spanning two consecutive days. Notably, 492 of these transactions are fraudulent, resulting in a significant class imbalance. The data is conveniently provided in a single file named creditcard.csv, which facilitates the seamless construction of an end-to-end fraud identification pipeline and a risk scoring task. At the field level, the dataset includes anonymized numerical features, commonly referred to as V1 to V28. It also includes variables that capture transaction time and amount, such as Time and Amount. A binary label is provided to distinguish between normal and fraudulent transactions. The dataset exhibits a clear temporal order and a strongly imbalanced structure. It is well-suited for method design that addresses two key challenges: rapid adaptation to limited samples and cross-scenario distribution shift. Techniques like time-based splits, subdomain sampling, or feature subspace construction can be employed to simulate the emergence of new fraud patterns and transfer across different scenarios. This aligns with the unified modeling objective that combines meta learning and domain adaptation.

B. Experimental Results

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

 doi:10.20944/preprints202602.0609.v1

Table 1. Comparative experimental results.

Method	Acc	Recall	Precision	F1-Score
DetectGNN[16]	0.742	0.721	0.754	0.737
FRAUD-RLA[17]	0.758	0.739	0.766	0.752
FinGraphFL[18]	0.764	0.746	0.771	0.758
CNN-GRU[19]	0.751	0.728	0.759	0.743
BiLSTM-Transformer[20]	0.769	0.752	0.777	0.764
Ours	0.781	0.763	0.789	0.775

The overall comparison shows that the proposed method achieves the best performance across all four metrics. Acc reaches 0.781. Recall is 0.763. Precision is 0.789. F1 is 0.775. Compared with existing methods, this consistent improvement indicates that the model is more robust in overall discrimination. It also achieves a more appropriate balance between minority class fraud detection and false alarm control. This matches the core requirements of stability and usability in a real open environment for credit fraud detection.

When comparing Ours with the strongest baseline, BiLSTM Transformer, Acc increases from 0.769 to 0.781. Recall increases from 0.752 to 0.763. Precision increases from 0.777 to 0.789. F1 increases from 0.764 to 0.775. More importantly, Recall and F1 rise together. This suggests that the gains are not achieved by simply shifting the threshold and sacrificing precision. Instead, the model reduces missed detections while maintaining a low false positive level. Under a shot and emerging fraud scenarios, this implies faster formation of an effective decision boundary from limited new patterns. It also reduces the risk exposure window during the cold start stage.

From the perspective of methodological differences, graph-based methods and sequence-based methods each have strengths but also clear limitations. Structured modeling approaches such as DetectGNN and FinGraphFL help capture relational fraud and coordinated group behavior. However, they can be sensitive to domain feature distribution changes during scenario transfer, which limits generalization. Temporal and attention-based architectures such as CNN, GRU, and BiLSTM Transformer are better at modeling sequential dependencies. Yet they can also suffer representation drift when new business channels appear or when data statistics change. Ours is more balanced across metrics. This indicates that jointly integrating meta learning and domain adaptation can mitigate both insufficient fast adaptation and cross-domain instability. It makes representations more transferable. It also makes updates more efficient. This is consistent with the rapid generalization objective of the paper.

The class weight coefficient directly alters the model's focus on minority class samples during training, thus affecting the learning preference for the discrimination boundary and the sensitivity of risk identification. To verify the impact of this hyperparameter on model behavior, independent observation and comparative analysis of a single indicator under different weight settings are required. The experimental results are shown in Figure 2.

Across the four subplots, the sensitivity patterns indicate that the class weight coefficient has a strong pulling effect on the decision boundary. This effect differs across evaluation metrics. This implies that in an open setting where few samples and emerging fraud coexist, the weight coefficient not only controls how much the model attends to minority class signals. It also changes representation stability and the direction of gradient updates. This can cause non-synchronized changes at the metric level. Because the method emphasizes fast adaptation and cross-domain robustness, the weight coefficient acts as a key control knob that links the two goals. Its reasonable range directly determines whether the model can maintain a consistent risk ranking in new scenarios.

 doi:10.20944/preprints202602.0609.v1

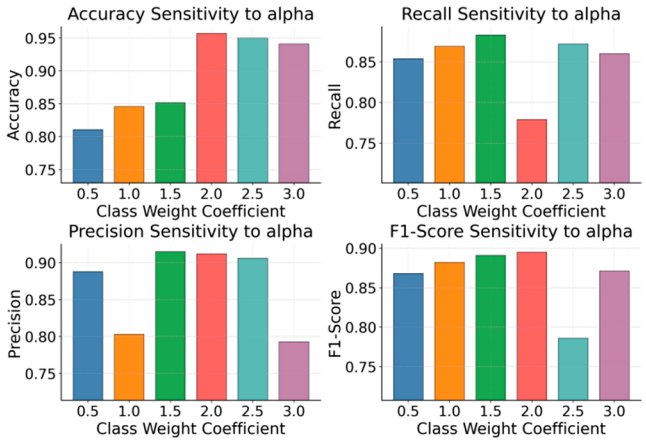

Figure 2. The impact of class weight coefficients on experimental results.

From the perspective of overall discrimination quality, increasing the weight from a low level first improves performance. The trend then becomes saturated and may slightly decline. This rise-then-flatten pattern usually indicates that moderately increasing the minority class contribution strengthens sensitivity to fraud patterns. It helps the model capture discriminative cues faster under limited labels. When the weight continues to grow, the training objective can become overly dominated by minority class gradients. This can weaken structural coverage and calibration for the majority class. It may further reduce stability during cross-domain transfer. For a framework that stresses domain adaptation, this suggests that weight tuning should not be treated as the only bias correction tool. More stable domain invariant representations should support the fast adaptation process.For Recall and Precision, the curves exhibit stronger oscillations and non-monotonic behavior. This indicates a clear trade-off and higher threshold sensitivity. A larger weight does not necessarily lead to a sustained increase in Recall. In some ranges, it can even trigger a structural reversal between missed detections and false alarms. This suggests that once attention to the minority class exceeds a reasonable level, the decision boundary may drift or become unstable. This aligns well with emerging fraud scenarios. New patterns often arrive simultaneously with a distribution shift and sample scarcity. If the weight setting does not match the target domain statistics, bias can be amplified, and generalization can deteriorate. Therefore, task-level fast adaptation and domain alignment are crucial to jointly buffer these discontinuous fluctuations.

The behavior of the aggregate metric further shows that there exists an effective range of class weights that better balances the costs of false positives and false negatives. Deviating from this range leads to a clear degradation in overall quality. Since the paper emphasizes rapid generalization under few-shot conditions, the goal is not to push the weight to an extreme to pursue a one-sided advantage. The goal is to keep inter-domain representations consistent while allowing meta learning updates to operate within a more stable loss geometry. This yields a more transferable optimum. In other words, these sensitivity results support treating the weight coefficient as a controllable hyperparameter for single metric inspection. They also support using domain adaptation to reduce its sensitivity to environmental change, so that the model maintains more consistent risk discrimination when transferring across time and scenarios.

V. Conclusion

This paper addresses the challenge of rapid generalization under few-shot and emerging fraud scenarios. It proposes a credit fraud detection framework that jointly integrates meta learning and domain adaptation. At the methodological level, it responds to two practical pain points in real operations. The first is insufficient learning caused by scarce labels during cold start in new scenarios. The second is performance instability caused by distribution shift across time, channels, and products. With shared representations and task-level fast adaptation, the model can adjust its

 doi:10.20944/preprints202602.0609.v1

decision boundary quickly with a small number of new samples. With representation alignment and distribution stability constraints, the model reduces reliance on a single historical distribution. It can therefore learn risk signals more continuously in dynamic environments. Overall, this study unifies fast adaptation and cross-domain robustness within one modeling paradigm. It offers a solution that better matches the evolutionary nature of open-world credit risk control.

From an application perspective, the framework has direct implications for financial institutions and platform-based businesses. It helps shorten the risk control preparation cycle when launching new services or integrating new channels. It also reduces monetary losses and chargeback costs caused by early-stage diffusion of emerging fraud. At the same time, it mitigates false positives and false negatives induced by distribution drift. This improves system usability and user experience. More importantly, the joint design reduces reliance on frequent full retraining and extensive manual hyperparameter tuning. It supports continuous operation through more transferable representations and more efficient adaptation mechanisms. This is particularly relevant for fast-changing scenarios such as high-frequency transactions, cross-border payments, installment credit, and buy-now-pay-later services. The same idea can be naturally extended to related tasks, including anti-money laundering, fake account detection, merchant risk rating, and transaction anomaly monitoring. It provides a methodological basis for building a more unified intelligent risk governance system.

Looking ahead, several directions deserve further investigation. First, finer-grained scenario characterization and task construction can be introduced into the meta learning process. Adaptation would then rely not only on a small number of labels but also on self-supervised signals and weak supervision feedback. This can strengthen early capture of emerging fraud variants. Second, at the domain adaptation level, online modeling of non-stationary drift and abrupt events can be further explored. The model could maintain stable calibration in continuously changing statistical environments. It could also reduce the risk caused by error accumulation from pseudo labels. Finally, with increasing requirements for data compliance and privacy protection, this framework can be combined with privacy-preserving computation, federated learning, or cross-institution collaboration mechanisms. This would enable multi-domain knowledge sharing and robust generalization without exposing sensitive data. It would also promote credit risk control from isolated optimization toward an intelligent protection system that generalizes across scenarios and entities, and evolves sustainably.

Reference

1. R. Bin Sulaiman, V. Schetinin and P. Sant, "Review of machine learning approach on credit card fraud detection," Human-Centric Intelligent Systems, vol. 2, no. 1, pp. 55-68, 2022.
2. F. K. Alarfaj, I. Malik, H. U. Khan, et al., "Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms," IEEE Access, vol. 10, pp. 39700-39715, 2022.
3. E. A. L. M. Btoush, X. Zhou, R. Gururajan, et al., "A systematic review of literature on credit card cyber fraud detection using machine and deep learning," PeerJ Computer Science, vol. 9, p. e1278, 2023.
4. G. Zioviris, K. Kolomvatsos and G. Stamoulis, "Credit card fraud detection using a deep learning multistage model," Journal of Supercomputing, vol. 78, no. 12, 2022.
5. M. Habibpour, H. Gharoun, M. Mehdipour, et al., "Uncertainty-aware credit card fraud detection using deep learning," Engineering Applications of Artificial Intelligence, vol. 123, p. 106248, 2023.
6. J. F. Roseline, G. Naidu, V. S. Pandi, et al., "Autonomous credit card fraud detection using machine learning approach," Computers and Electrical Engineering, vol. 102, p. 108132, 2022.
7. A. Cherif, A. Badhib, H. Ammar, et al., "Credit card fraud detection in the era of disruptive technologies: A systematic review," Journal of King Saud University-Computer and Information Sciences, vol. 35, no. 1, pp. 145-174, 2023.
8. Z. Xu, X. Liu, Q. Xu, X. Su, X. Guo and Y. Wang, "Time series forecasting with attention-augmented recurrent networks: A financial market application," Proceedings of the 2025 2nd International Conference on Computer and Multimedia Technology, pp. 155-159, 2025.

 doi:10.20944/preprints202602.0609.v1

9. S. Li, Y. Wang, Y. Xing and M. Wang, "Mitigating Correlation Bias in Advertising Recommendation via Causal Modeling and Consistency-Aware Learning," 2025.
10. J. Lai, A. Xie, H. Feng, Y. Wang and R. Fang, "Self-Supervised Learning for Financial Statement Fraud Detection with Limited and Imbalanced Data," 2025.
11. Y. Wang, R. Fang, A. Xie, H. Feng and J. Lai, "Dynamic Anomaly Identification in Accounting Transactions via Multi-Head Self-Attention Networks," arXiv preprint arXiv:2511.12122, 2025.
12. H. Wang, C. Nie and C. Chiang, "Attention-Driven Deep Learning Framework for Intelligent Anomaly Detection in ETL Processes," 2025.
13. J. Li, Q. Gan, R. Wu, C. Chen, R. Fang and J. Lai, "Causal Representation Learning for Robust and Interpretable Audit Risk Identification in Financial Systems," 2025.
14. Oztas, B., Cetinkaya, D., Adedoyin, F., Budka, M., Aksu, G., and Dogan, H., "Transaction monitoring in anti-money laundering: a qualitative analysis and points of view from industry," Future Generation Computer Systems, vol. 159, pp. 161-171, 2024.
15. J. Zheng, H. Zhang, X. Yan, R. Hao and C. Peng, "Contrastive Knowledge Transfer and Robust Optimization for Secure Alignment of Large Language Models," arXiv preprint arXiv:2510.27077, 2025.
16. I. Sultana, S. M. Maheen, N. Kshetri, et al., "detectGNN: Harnessing Graph Neural Networks for Enhanced Fraud Detection in Credit Card Transactions," Proceedings of the 2025 13th International Symposium on Digital Forensics and Security, pp. 1-6, 2025.
17. D. Lunghi, Y. Molinghen, A. Simitsis, et al., "FRAUD-RLA: A new reinforcement learning adversarial attack against credit card fraud detection," arXiv preprint arXiv:2502.02290, 2025.
18. Z. Xia and S. C. Saha, "FinGraphFL: Financial Graph-Based Federated Learning for Enhanced Credit Card Fraud Detection," Mathematics, vol. 13, no. 9, p. 1396, 2025.
19. A. Chidananda, "Credit Card Fraud Detection Using Hybrid Deep Learning CNN-LSTM and CNN-GRU Models," M.S. thesis, California State University, Northridge, 2025.
20. P. Feng, "Hybrid BiLSTM-Transformer Model for Identifying Fraudulent Transactions in Financial Systems," Journal of Computer Science and Software Applications, vol. 5, no. 3, 2025.