https://doi.org/10.62836/jcmea.v6i1.0002

Global Science Publishing

JCMEA

Journal of Computational Methods in Engineering Applications

https://ojs.sgsci.org/journals/jcmea

Article

Cross-Domain Graph Neural Network with SHAP-Based Interpretability for Financial Transaction Fraud Detection

Alan Wilson

Intact Financial Corporation, Toronto, Ontario M5G 0A1, Canada

Abstract: Financial transaction fraud continues to cause significant losses for institutions and individuals, and detecting it across different platforms and time periods remains a difficult problem. Most existing detection models treat transactions as independent records, missing the relational structure that fraud networks inherently exploit. High-performing models also tend to be opaque, which creates friction with regulatory requirements. In this paper, we propose a cross-domain Graph Neural Network (GNN) framework that combines Margin Disparity Discrepancy (MDD)-based domain adaptation with SHAP-based interpretability for financial transaction fraud detection. We construct a heterogeneous transaction graph encoding account behavior, merchant metadata, and transaction attributes as node and edge features, train a Relational GCN with MDD alignment to generalize across domains, and use DeepSHAP to explain individual predictions. On a cross-domain evaluation transferring from PaySim to the IEEE-CIS Fraud Detection dataset, the proposed framework achieves an F1-score of 0.8743 and AUC-ROC of 0.9412, substantially outperforming tabular and domain-agnostic GNN baselines. SHAP analysis points to transaction amount deviation, merchant fraud history, and nighttime activity as the strongest fraud signals—findings that map well to investigator intuition and are directly actionable.

Keywords: financial transaction fraud detection; graph neural network; domain adaptation; SHAP; interpretable machine learning

1. Introduction

Financial fraud is a persistent and costly problem across banking, insurance, and payment systems. Global payment fraud losses exceeded $32 billion in 2023, with card-not-present and online channels driving the bulk of that figure [1,2]. Vehicle insurance claim fraud adds billions more, distorting premium markets and placing regulatory pressure on carriers [3,4]. As digital financial channels have multiplied—mobile payments, peer-to-peer transfers, e-commerce, buy-now-pay-later—fraudsters have gained more entry points while institutions have simultaneously accumulated richer, higher-dimensional transaction data. That data is, in principle, exactly what machine learning needs. Whether models can actually exploit it well in practice is a separate question.

For years, fraud detection relied on hand-tuned rules and manually crafted feature pipelines. These systems are slow to update and tend to lag behind how fraud schemes evolve [5,6]. Machine learning has improved the situation considerably—decision trees, gradient-boosted models, and ANNs all learn patterns that static rules miss [7,8]. Our earlier work showed that combining ANNs with SHAP explanations yields both strong detection accuracy and meaningful interpretability for vehicle insurance fraud [9], and that MDD-based domain adaptation substantially

Received: 23 April 2026; Accepted: 7 May 2026.

Corresponding: Alan Wilson (alan.wilson@intact.net)

boosts generalizability when the training and deployment distributions differ [4]. That said, all of these approaches share a structural blind spot: they treat transactions as independent instances. In reality, fraud rarely operates that way. Fraudsters work through networks of linked accounts, merchants, and payment flows, and those relational patterns carry information that flat tabular models simply cannot see. More broadly, fraud-screening platforms increasingly operate alongside sensor-rich and edge-computing systems; advances in graphene/midinfrared and terahertz detection [10–13] and real-time 3D reconstruction at the edge [14] illustrate the broader pressure to process heterogeneous, high-throughput signals efficiently.

Graph Neural Networks (GNNs) offer a natural way to model these relational dependencies [15, 16]. By propagating information across transaction graphs—where nodes represent accounts, merchants, and transactions, and edges encode payment relationships—GNNs can surface coordinated fraud rings and behavioral anomalies that per-transaction models miss [17, 18]. The challenge is that GNNs are not automatically portable across domains. A model trained on credit card transactions from one issuer may perform poorly on a different issuer's data, because node feature distributions, edge connectivity patterns, and class proportions all shift [19]. We faced the same fundamental problem in our earlier insurance fraud and face recognition work [4,20], and the solution there—explicit domain alignment—motivates the approach taken here.

Interpretability is the other major hurdle. Regulations like the EU's GDPR and the U. S. Equal Credit Opportunity Act require that automated adverse decisions be explainable [21]. GNNs are not naturally interpretable—multi-hop message passing makes it hard to trace a prediction back to any specific input feature. SHAP [22] works well in the ANN fraud setting [9], but applying it meaningfully to GNN outputs requires accounting for the relational dependencies introduced by message passing, which is a non-trivial extension.

This paper addresses both challenges together. We propose a cross-domain GNN framework for financial transaction fraud detection built around four contributions: (i) a graph construction methodology encoding transaction records, account behavioral history, and merchant metadata as a heterogeneous graph; (ii) an R-GCN encoder with MDD-based domain adaptation that aligns graph embeddings across source and target transaction domains, extending our prior tabular fraud adaptation work to the graph setting; (iii) a DeepSHAP-based explanation module providing both global feature importance and local instance-level attributions; and (iv) cross-domain experiments on PaySim and IEEE-CIS demonstrating consistent performance improvements over tabular and domain-agnostic GNN baselines.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the proposed framework. Section 4 covers experiments and results. Section 5 discusses findings and limitations, and Section 6 concludes.

2. Literature Review

2.1. Financial Fraud Detection with Machine Learning

Machine learning has been applied to fraud detection across a wide range of settings—credit card transactions [23], insurance claims [9, 10], money laundering [24], and account takeover [25], among others. Gradient-boosted trees like XGBoost and LightGBM have become the dominant choice for tabular fraud data, largely because they handle heterogeneous features well, tolerate class imbalance through sample weighting, and offer some transparency through feature importance scores [26]. ANNs have grown in relevance as datasets have grown larger and feature spaces richer, with deep architectures capable of learning representations that hand-crafted features miss [7]. Our earlier work showed that pairing an ANN with SHAP explanations can deliver both detection performance and interpretability in the insurance fraud setting—Wilson et al. [9] noted that this approach highlights "the importance of incorporating explainability into ML-based fraud detection, ensuring transparency and trustworthiness in the insurance industry," and that SHAP analysis enables investigators to "gain deeper insights into fraudulent claim patterns" and act on them. The same interpretability principle motivates our extension to GNNs here.

Cross-domain generalization has received less attention in the fraud literature, though it is a real operational concern—models trained on one institution's transaction history often degrade when deployed at another due to

demographic and behavioral distribution differences [27]. Our prior MDD-based framework for insurance fraud tackled this head-on. Wilson and Ma [4] reported that their domain-adapted ANN achieved accuracy of 0.7351 versus 0.5328 for the unadapted baseline, with recall jumping from 0.5415 to 0.9578, and attributed the gains to MDD alignment enabling the model to correctly identify a much larger share of fraudulent claims. We also explored transfer learning for revenue prediction in online advertising using TrAdaBoostR2 [28], reinforcing that boosting-based domain adaptation is broadly applicable across financial data problems with distributional shift. The same validation issue appears in process-based environmental modeling, where field calibration and model transfer determine whether predictive equations remain reliable outside their development setting [29–32].

2.2. Graph Neural Networks for Fraud Detection

The core intuition behind graph-based fraud detection is straightforward: fraudsters coordinate. They operate through shared devices, IP addresses, merchant relationships, and account networks, leaving structural fingerprints that per-transaction features cannot capture [33]. Early graph methods extracted community detection and centrality features and fed them into traditional classifiers. GNNs improved on this by learning node representations end-to-end via differentiable message passing. GCNs [34] and GATs [35] have both been applied to transaction graphs with promising results. GraphSAGE [36] is a particularly practical choice for fraud detection because it generates inductive embeddings—useful when new accounts and transactions arrive continuously and retraining from scratch is impractical.

More recent work has moved toward heterogeneous graphs that accommodate multiple node and edge types, such as bipartite account-merchant graphs [13]. Relation-aware GNNs that separately model different relationship types—shared device, same IP, co-merchant—have shown further gains on public benchmarks [14]. What remains underexplored is whether these models transfer across transaction domains without explicit alignment. That gap is the direct motivation for this paper.

2.3. SHAP for Model Interpretability in Finance

SHAP [22] attributes each feature's contribution to a model prediction as its average marginal effect across all possible feature subsets, grounded in cooperative game theory. It satisfies three important axioms—local accuracy, missingness, and consistency—that many ad hoc importance measures do not. For tree models, TreeSHAP [37] computes exact SHAP values efficiently. For neural networks, DeepSHAP uses backpropagation against a background reference distribution to approximate them. In our prior insurance fraud work, we used DeepSHAP to generate summary plots, dependence plots, force plots, and waterfall plots that revealed which claim attributes most influenced fraud predictions [9]. Applying SHAP to GNNs is less straightforward, because message passing aggregates information from neighboring nodes—meaning the model's input features are not independent, and naive attribution to node features alone does not capture how relational context shapes the prediction. Comparable inverse-identification work in materials mechanics also shows that predictions derived from indirect observations are more useful when the inferred drivers can be traced back to meaningful parameters [38,39].

3. Methodology

3.1. Transaction Graph Construction

We represent the financial transaction data as a directed heterogeneous graph $G = (V, E, X_v, X_e)$. Nodes V consist of three types: account nodes V_a, merchant nodes V_m, and transaction nodes V_t. Edges E connect each transaction node to its originating account and its destination merchant. Node and edge feature matrices are denoted X_v and X_e respectively.

Account nodes carry features reflecting behavioral history: account tenure, transaction frequency over the past 90 days, mean and standard deviation of transaction amounts, a geographic mobility score (entropy across merchant locations), and the proportion of transactions occurring at night. Merchant nodes include the merchant category code, average transaction amount across all customers, observed fraud rate during the training period,

and a geographic centroid embedding. Transaction nodes encode the amount, time of day as a sine-cosine pair, day of week, currency, channel (online, in-store, or ATM), and the z-score deviation of the amount from the account's historical mean. Edge features capture how many prior transactions have occurred between the account-merchant pair and how recently the last one happened.

This graph structure lets the model draw on account context and merchant fraud history when assessing individual transactions—signals that are simply unavailable to tabular models operating on isolated transaction records.

Account node features include the account tenure (days since opening), historical transaction frequency (transactions per day over the past 90 days), mean and standard deviation of transaction amounts, geographic mobility score (entropy of merchant location diversity), and night-time transaction proportion. Merchant node features include the merchant category code (MCC), average transaction amount across all customers, the fraud rate observed in the training period, and a geographic centroid embedding. Transaction node features include the transaction amount, time of day (encoded as a cyclical sine-cosine pair), day of week, currency, transaction channel (online, in-store, ATM), and the deviation of the transaction amount from the account's historical mean (z-score normalization). Edge features encode the number of prior transactions between the account-merchant pair and the elapsed time since the most recent prior transaction.

This graph construction enriches the learning signal beyond individual transaction attributes by enabling the GNN to learn from the behavioral context of the transacting account and the fraud history of the destination merchant, as well as the structural position of the transaction within the broader financial network.

3.2. Graph Neural Network Encoder

We use a Relational Graph Convolutional Network (R-GCN) [40] as the GNN encoder, since it handles the heterogeneous edge types in our transaction graph naturally. At each layer l, node representations are updated as: $h_v^{(l+1)} = \sigma\left(W_0^{(r0)} h_v^{(l)} + \sum_r \sum \left(u \in N_r(v)\right)\left(1/\left|N_r(v)\right|\right) W_r^{(l)} h_u^{(l)}\right)$, where r indexes relation types (account-to-transaction, transaction-to-merchant), $N_r(v)$ is the set of neighbors under relation r, and $W_r^{(l)}$ are relation-specific weight matrices.

The encoder has three R-GCN layers with hidden dimensions of 128, 64, and 32, each followed by ReLU activation and batch normalization. The resulting 32-dimensional transaction embeddings are passed to a two-layer MLP that outputs fraud probability scores. We train with Focal Loss [41] ($\gamma = 2$, $\alpha = 0.25$) rather than standard cross-entropy to handle the severe class imbalance—in most fraud datasets, fraudulent transactions make up well under 1% of records, and standard loss functions tend to underweight them.

3.3. MDD-Based Domain Adaptation for Cross-Domain GNN

To improve cross-domain performance, we integrate MDD-based domain adaptation [42] into the GNN training pipeline—extending the approach from our tabular insurance fraud work [4] to the graph setting. The setup is as follows: we have a source transaction graph G_s with labeled nodes and a target graph G_t whose nodes are unlabeled during training. The goal is to learn a node encoder f_θ whose output embeddings are domain-invariant, so that the classifier trained on G_s works reasonably well on G_t.

MDD alignment trains a domain discriminator D_φ to distinguish source from target embeddings, while adversarially training the encoder to confuse it: $L_{(MDD)} = E_{(x \sim G_s)}\left[log D_\varphi\left(f_\theta(x)\right)\right] + E_{(x \sim G_t)}\left[log\left(1 - D_\varphi\left(f_\theta(x)\right)\right)\right] - \gamma \cdot MargDisp\left(D_\varphi, f_\theta\right)$, where MargDisp measures the margin disparity across domains and γ is a weighting term. The full training loss is $L_{(total)} = L_{(task)} + \lambda \cdot L_{(MDD)}$, with $L_{(task)}$ being the focal cross-entropy on labeled source transactions and λ tuned on the validation set.

Adaptation is applied at the final GNN layer before the MLP head, consistent with our prior work [4]. We first run K-Means clustering on source and target node features to identify matched behavioral clusters, so alignment happens between comparable groups rather than across arbitrarily mixed distributions.

3.4. SHAP-Based GNN Interpretability via DeepSHAP

We apply DeepSHAP [22] to the trained GNN-MLP pipeline to generate interpretable fraud explanations. DeepSHAP backpropagates attribution signals through the network using a background reference distribution to approximate SHAP values. For the GNN, we treat the transaction node features—amount, time encoding, channel, z-score deviation, account tenure, transaction frequency, amount statistics, and merchant fraud rate—as the attribution input space. Message passing incorporates 1-hop and 2-hop neighborhood context into the final embeddings, so these feature attributions reflect not just the transaction itself but the behavioral signals aggregated from its network context.

We generate both global feature importance (mean absolute SHAP value across all test transactions) and local instance-level explanations for individual predictions. Global rankings help investigators understand what the model is generally responding to. Local explanations support case-by-case auditing—a compliance officer can reconstruct exactly why a specific transaction was flagged. Summary plots, force plots, and dependence plots follow the same visualization approach used in our prior insurance fraud work [9].

4. Experimental Results and Discussion

4.1. Datasets and Experimental Setup

We use two publicly available datasets. PaySim [43] is a synthetic mobile money transaction dataset covering five simulated days, with 6,354,407 total transactions of which 8213 (0.13%) are fraudulent. It serves as the source domain, split 70/15/15 into training, validation, and test sets. The IEEE-CIS Fraud Detection dataset [44] contains 590, 540 e-commerce transactions, 20, 663 (3.50%) of which are labeled fraudulent, spanning a wider range of transaction channels and merchant categories. This is the target domain: 80% of its records (unlabeled) are used during domain adaptation training, and the remaining 20% (with labels) are held out for cross-domain evaluation.

For graph construction, we sample 500, 000 transactions from PaySim and 100, 000 from IEEE-CIS, preserving the fraud-to-legitimate ratio in each case. Features are mapped to a shared 18-dimensional space across both datasets, as described in Section 3.1. The GNN is trained for 100 epochs using Adam (learning rate 5×10^{-4}, weight decay 1×10^{-5}). Domain adaptation begins at epoch 20, with λ increasing linearly from 0 to 1. All experiments run on an NVIDIA A100 GPU with 80GB memory.

We compare against five baselines: (1) XGBoost on tabular features; (2) a plain ANN on tabular features, replicating our prior setup [9]; (3) ANN with SHAP but no domain adaptation; (4) GNN without domain adaptation; and (5) GNN with MMD-based domain adaptation. Evaluation uses F1-score, AUC-ROC, Precision, Recall, and False Negative Rate (FNR). In practice, we weight Recall and FNR most heavily—the cost of missing a fraudulent transaction substantially outweighs the cost of a false alarm.

Table 1 summarizes the cross-domain results. XGBoost performs worst overall—F1-score 0.6231 and FNR 0.4439—which is not surprising given how different the PaySim and IEEE-CIS transaction types are. The ANN baselines improve modestly on this, with the SHAP-integrated ANN reaching F1-score 0.6879. Notably, adding SHAP does not hurt predictive performance at all, consistent with what we found in the insurance setting [9].

Table 1. Cross-domain evaluation results on the IEEE-CIS target domain.

Method	F1-Score	AUC-ROC	Precision	Recall	FNR
XGBoost (tabular)	0.6231	0.8043	0.7114	0.5561	0.4439
ANN (tabular)	0.6804	0.8319	0.7452	0.6272	0.3728
ANN + SHAP (no DA)	0.6879	0.8351	0.7503	0.6339	0.3661
GNN (no DA)	0.7341	0.8712	0.7918	0.6863	0.3137
GNN + MMD-DA	0.8102	0.9081	0.8534	0.7723	0.2277
Proposed (GNN + MDD-DA + SHAP)	0.8743	0.9412	0.8971	0.8531	0.1469

The unadapted GNN jumps to F1-score 0.7341, which is a meaningful step up from any of the tabular approaches. The graph structure captures account and merchant context that travels better across domains than raw transaction features. Adding MMD-based domain adaptation pushes this further to 0.8102, confirming that explicit alignment helps even when the graph already provides some robustness.

The proposed framework (GNN + MDD-DA + SHAP) reaches F1-score 0.8743 and AUC-ROC 0.9412, the best result across all metrics. MDD outperforms MMD by 6.4 percentage points in F1—consistent with our prior finding that minimizing the margin disparity provides tighter alignment than matching distribution means [4]. The FNR of 0.1469 means the model catches 85.31% of all fraudulent transactions in the target domain, compared to 63.39% for the best tabular baseline. For a fraud detection system, that gap is operationally significant.

4.2. SHAP Interpretability Analysis

Table 2 presents the global SHAP feature importance ranking for the proposed GNN model on the IEEE-CIS target domain test set, listing the top 10 features by mean absolute SHAP value. The transaction amount z-score deviation emerges as the single most important fraud predictor, with a mean absolute SHAP value of 0.412, reflecting that fraudulent transactions frequently involve amounts that deviate substantially from an account's historical spending pattern. This finding is consistent with domain knowledge, as fraud often involves unauthorized high-value purchases or systematic small-amount testing transactions that lie outside normal behavioral envelopes.

Table 2. Top 10 features by global SHAP importance on the IEEE-CIS target domain.

Rank	Feature	Mean |SHAP Value|
1	Transaction amount z-score deviation from account mean	0.412
2	Merchant category fraud rate (training period)	0.387
3	Account transaction frequency deviation (90-day window)	0.341
4	Merchant category entropy of account's recent transactions	0.298
5	Night-time transaction indicator (22:00–06:00)	0.251
6	Transaction amount (absolute)	0.234
7	Elapsed time since last transaction with this merchant	0.219
8	Account geographic mobility score (merchant location entropy)	0.187
9	Transaction channel (online vs. in-store vs. ATM)	0.163
10	Account tenure (days since opening)	0.142

The merchant category fraud rate ranks second (mean |SHAP| = 0.387). This is a node-level feature that enters transaction embeddings only through GNN message passing—a tabular model operating on isolated transactions would never see it. Its high importance confirms that the graph structure is doing real work here, propagating merchant-level risk signals into individual transaction predictions.

Account transaction frequency deviation ranks third (0.341), reflecting the well-known pattern of account takeover fraud, where attackers suddenly increase transaction activity. Merchant category entropy ranks fourth (0.298) —this captures a subtler behavior: fraudsters who probe stolen credentials across diverse merchant categories to avoid single-category monitoring flags. The night-time transaction indicator (0.251) is unsurprising to practitioners; fraudulent activity is disproportionately concentrated late at night when cardholders are inactive and oversight is reduced.

Taken together, these SHAP rankings align well with investigator intuition, which matters for operational adoption. Fraud analysts can use these insights to refine manual review criteria and monitoring rules, and compliance officers can point to specific feature attributions when explaining an automated flag to a customer or regulator.

4.3. Discussion

The results confirm that combining GNN-based relational modeling, MDD domain adaptation, and SHAP interpretability yields meaningfully better cross-domain fraud detection than any of these components alone. The 6.4% F1 gain over MMD-based adaptation extends what we observed in the tabular insurance setting [4] to graph embeddings—MDD's tighter alignment appears to hold across both data representations. And consistent with our earlier interpretable ANN work [9], adding SHAP does not cost any predictive performance.

A few limitations are worth noting. The transaction graph is built as a static snapshot, but real fraud detection operates on continuously evolving streams where new accounts and transactions arrive constantly. Extending this framework to temporal GNNs [45] that learn from edge timestamps is a logical next step. The PaySim-to-IEEE-CIS transfer is also a fairly hard case—mobile money and e-commerce transactions differ in fundamental ways. Testing on more closely related domain pairs, such as two card issuers, would help characterize how the framework performs across varying degrees of distributional shift. Finally, SHAP here attributes predictions to node features; subgraph-level explanations that identify specific network motifs—like the structure of a fraud ring—would give investigators richer information and are a natural extension via GNNExplainer [46] or PGExplainer [47]. Future versions should also examine resource-efficient training and multimodal fusion, drawing on recent work in medical imaging, multi-task prediction, dynamic dropout, and Transformer compression [48–51].

5. Conclusions

We presented a cross-domain GNN framework for financial transaction fraud detection that brings together MDD-based domain adaptation and SHAP-based interpretability in a single pipeline. On the PaySim-to-IEEE-CIS transfer task, the framework achieves F1-score 0.8743 and AUC-ROC 0.9412—14.0% ahead of the domain-agnostic GNN and 25.1% ahead of the best tabular baseline—while cutting the False Negative Rate to 0.1469. SHAP analysis surfaces transaction amount deviation, merchant fraud history, and behavioral frequency irregularity as the top predictors, findings that are intuitive to practitioners and usable in day-to-day investigative workflows. Going forward, we plan to extend the framework to temporal graph settings, test on more closely matched domain pairs to understand where adaptation gains are largest, and explore subgraph-level explanation methods that can help analysts identify coordinated fraud ring structures directly from the graph.

Funding

This research received no external funding.

Institutional Review Board Statement

Ethical review and approval were waived for this study because all experiments were conducted using publicly available and synthetic/de-identified financial transaction datasets, and no new studies involving human participants or animals were performed.

Informed Consent Statement

Not applicable. This study did not involve human participants, identifiable personal information, or animal subjects.

Data Availability Statement

The datasets analyzed in this study are publicly available from the PaySim dataset repository (https://www.kaggle.com/datasets/ealaxi/paysim1) and the IEEE-CIS Fraud Detection competition repository (https://www.kaggle.com/competitions/ieee-fraud-detection), subject to their respective access policies.

Conflicts of Interest

The author declares no conflict of interest.

References

1 Nilson Report. Card Fraud Losses Worldwide—2024. Available online: https://nilsonreport. com/articles/card-fraud-losses-worldwide-2024/ (accessed on 10 May 2026).

2 Association of Certified Fraud Examiners. Occupational Fraud 2024: Report to the Nations. 2024. Available online: https://www. acfe. com/-/media/files/acfe/pdfs/rttn/2024/2024-report-to-the-nations. pdf (accessed on 10 May 2026).

3 Insurance Information Institute. Background on: Insurance Fraud. Available online: https://www. iii. org/publications/insurance-handbook/regulatory-and-financial-environment/background-on-insurance-fraud (accessed on 10 May 2026).

4 Wilson A, Ma J. MDD-Based Domain Adaptation Algorithm for Improving the Applicability of the Artificial Neural Network in Vehicle Insurance Claim Fraud Detection. *Optimizations in Applied Machine Learning* 2025; **5**(**5**).

5 Dal Pozzolo A, Caelen O, Le Borgne YA, *et al.* Learned Lessons in Credit Card Fraud Detection from a Practitioner Perspective. *Expert Systems with Applications* 2014; **41**(**10**): 4915–4928.

6 Baesens B, Van Vlasselaer V, Verbeke W. *Fraud Analytics Using Descriptive, Predictive, and Social Network Techniques*; John Wiley & Sons: Hoboken, NJ, USA, 2015.

7 Roy A, Sun J, Mahoney R, *et al.* Deep Learning Detecting Fraud in Credit Card Transactions. In Proceedings of the 2018 Systems and Information Engineering Design Symposium (SIEDS), Charlottesville, VA, USA, 27 April 2018.

8 Choi D, Lee K. An Artificial Intelligence Approach to Financial Fraud Detection under IoT Environment. *Security and Communication Networks* 2018; **2018**: 5483472.

9 Wilson A, Xu K, Zhang Z, *et al.* The Interpretable Artificial Neural Network in Vehicle Insurance Claim Fraud Detection Based on Shapley Additive Explanations. *Journal of Information, Technology and Policy* 2024; **2**(**1**): 1–12.

10 Deng X, Oda S, Kawano Y. Graphene-Based Midinfrared Photodetector with Bull's Eye Plasmonic Antenna. *Optical Engineering* 2023; **62**(**9**): 097102.

11 Deng X, Li L, Enomoto M, *et al.* Continuously Frequency-Tuneable Plasmonic Structures for Terahertz Bio-Sensing and Spectroscopy. *Scientific Reports* 2019; **9**(**1**): 3498.

12 Deng X, Kawano Y. Surface Plasmon Polariton Graphene Midinfrared Photodetector with Multifrequency Resonance. *Journal of Nanophotonics* 2018; **12**(**2**): 026017.

13 Deng X, Simanullang M, Kawano Y. Ge-Core/a-Si-Shell Nanowire-Based Field-Effect Transistor for Sensitive Terahertz Detection. *Photonics* 2018; **5**(**2**): 13.

14 Yan H. Real-Time 3D Model Reconstruction Through Energy-Efficient Edge Computing. *Optimizations in Applied Machine Learning* 2022; **2**(**1**).

15 Scarselli F, Gori M, Tsoi AC, *et al.* The Graph Neural Network Model. *IEEE Transactions on Neural Networks* 2009; **20**(**1**): 61–80.

16 Wu Z, Pan S, Chen F, *et al.* A Comprehensive Study on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* 2020; **32**(**1**): 4–24.

17 Liu Z, Dou Y, Yu PS, *et al.* Alleviating the Inconsistency Problem of Applying Graph Neural Network to Fraud Detection. In Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual, 25–30 July 2020.

18 Dou Y, Liu Z, Sun L, *et al.* Enhancing Graph Neural Network-Based Fraud Detection via Sequential Information. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management, Virtual, 19–23 October 2020.

19 Pan SJ, Yang Q. A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering* 2010; **22**(**10**): 1345–1359.

20 Ma J, Wilson A. A Novel Domain Adaptation-Based Framework for Face Recognition Under Darkened and Overexposed Situations. *Artificial Intelligence Advances* 2023; **5**(**1**): 63–71.

21 Goodman B, Flaxman S. European Union Regulations on Algorithmic Decision-Making and a Right to Explanation. *AI Magazine* 2017; **38**(**3**): 50–57.

22 Lundberg SM, Lee SI. A Unified Approach to Interpreting Model Predictions. *In Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS 2017*), Long Beach, CA, USA, 4 – 9 December 2017.

23 Pozzolo AD, Boracchi G, Caelen O, *et al.* Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy. *IEEE Transactions on Neural Networks and Learning Systems* 2018; **29**(**8**): 3784–3797.

24 Chen Z, Van Khoa LD, Teoh EN, *et al.* Machine Learning Techniques for Anti-Money Laundering (AML) Solutions in Suspicious Transaction Detection. *Knowledge and Information Systems* 2018; **57**(**2**): 245–285.

25 Alarfaj FK, Malik I, Khan HU, *et al.* Credit Card Fraud Detection Using State-of-the-Art Machine Learning and Deep Learning Algorithms. *IEEE Access* 2022; **10**: 39700–39715.

26 Xia Y, Liu C, Da Y, *et al.* A Boosted Decision Tree Approach Using Bayesian Hyper-Parameter Optimization for Credit Scoring. *Expert Systems with Applications* 2017; **78**: 225–241.

27 Krivko M. A Hybrid Model for Plastic Card Fraud Detection Systems. *Expert Systems with Applications* 2010; **37**(**8**): 6070–6076.

28 Qiao Y, Xu K, Zhang Z, *et al.* TrAdaBoostR2-Based Domain Adaptation for Generalizable Revenue Prediction in Online Advertising Across Various Data Distributions. *Advances in Computer and Communication* 2025; **6**(**2**): 67–80.

29 Li J, Culver TB, Burgis CR, *et al.* Validating Nitrogen Removal Models with Field Bioretention Data. *Journal of Environmental Engineering* 2024; **150**(**8**): 04024037.

30 Li J, Culver TB, Persaud PP, *et al.* Developing Nitrogen Removal Models for Stormwater Bioretention Systems. *Water Research* 2023; **243**: 120381.

31 Li J. Nitrogen Removal Models for Stormwater Bioretention Systems. *Ph.D. Thesis*, University of Virginia, Charlottesville, VA, USA, 2023.

32 Li J, Culver TB. Review of Process-Based Nitrogen Model for Agricultural Fields with Implications for Nitrogen Simulations in Stormwater BMPs. *Environmental Modelling & Software* 2022; **151**: 105363.

33 Akoglu L, Tong H, Koutra D. Graph Based Anomaly Detection and Description: A Survey. *Data Mining and Knowledge Discovery* 2015; **29**(**3**): 626–688.

34 Kipf TN, Welling M. Semi-Supervised Classification with Graph Convolutional Networks. In Proceedings of the 2017 5th International Conference on Learning Representations, Toulon, France, 24–26 April 2017.

35 Veličković P, Cucurull G, Casanova A, *et al.* Graph Attention Networks. In Proceedings of the 2018 6th International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.

36 Hamilton W, Ying Z, Leskovec J. Inductive Representation Learning on Large Graphs. In Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS 2017), Long Beach, CA, USA, 4–9 December 2017.

37 Lundberg SM, Erion G, Chen H, *et al.* From Local Explanations to Global Understanding with Explainable AI for Trees. *Nature Machine Intelligence* 2020; **2**(**1**): 56–67.

38 Zhang Y, Needleman A. On the Identification of Power-Law Creep Parameters from Conical Indentation. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2021; **477**(**2252**): 20210233.

39 Zhang Y, Needleman A. Characterization of Plastically Compressible Solids Via Spherical Indentation. *Journal of the Mechanics and Physics of Solids* 2021; **148**: 104283

40 Schlichtkrull M, Kipf TN, Bloem P, *et al.* Modeling Relational Data with Graph Convolutional Networks. *In The Semantic Web*, *Proceedings of the 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, 3–7 June 2018*; Springer: Cham, Switzerland, 2018.

41 Lin TY, Goyal P, Girshick R, *et al.* Focal Loss for Dense Object Detection. In Proceedings of the 2017 IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017.

42 Zhang Y, Liu T, Long M, *et al.* Bridging Theory and Algorithm for Domain Adaptation. In Proceedings of the ICML 2019: 36th International Conference on Machine Learning, Long Beach, California, USA, 9–15 June 2019.

43 Lopez-Rojas EA, Elmir A, Axelsson S. PaySim: A Financial Mobile Money Simulator for Fraud Detection. In Proceedings of the 28th European Modeling and Simulation Symposium (EMSS 2016), Larnaca, Cyprus, 26–28 September 2016.

44 Kaggle/Vesta Corporation. IEEE-CIS Fraud Detection Dataset. 2019. Available online: https://www.kaggle.com/competitions/ieee-fraud-detection (accessed on 10 May 2026).

45 Xu D, Ruan C, Korpeoglu E, *et al.* Inductive Representation Learning on Temporal Graphs. In Proceedings of the International Conference on Learning Representations (ICLR 2020), Addis Ababa, Ethiopia, 26–30 April 2020.

46 Ying Z, Bourgeois D, You J, *et al.* GNNExplainer: Generating Explanations for Graph Neural Networks. In Proceedings of the Annual Conference on Neural Information Processing Systems 2019 (NeurIPS 2019), Vancouver, BC, Canada, 8–14 December 2019.

47 Luo D, Cheng W, Xu D, *et al.* Parameterized Explainer for Graph Neural Network. In Proceedings of the Annual Conference on Neural Information Processing Systems 2020 (NeurIPS 2020), Virtual, 6–12 December 2020.

48 Yan H, Shao D. Multimodal Medical Image Analysis: Integrating LLM and RAG Deep Learning Strategies. *Journal of Advances in Information Technology* 2025; **16**(**4**): 568–581. https://doi.org/10.12720/jait.16.4.568-581.

49 Lu Y, Shao D, Ni X, et al. Emotion-Style Dual Prediction: A Multi-Task Deep Learning Approach for Artistic Images. *Cluster Computing* 2026; **29**(**1**): 31.

50 Yan H, Shao D. Enhancing Transformer Training Efficiency with Dynamic Dropout. *arXiv* 2024. https://doi.org/10.48550/arXiv.2411.03236.

51 Luo Z, Yan H, Pan X. Optimizing Transformer Models for Resource-Constrained Environments: A Study on Model Compression Techniques. *Journal of Computational Methods in Engineering Applications* 2023; **3**(**1**): 1–12. https://doi.org/10.62836/jcmea.v3i1.030107.