

Deep Learning Approaches for Intelligent Fraud Detection in Digital Banking Transactions

1st Tarun Vakkalagadda

Data Scientist - Intern

tarun.vakkalagaddaa@gmail.com

Abstract—The rapid development of digital banking has not only enhanced convenience for users but also amplified the challenge of protecting against fraud, which remains both expensive and harmful for banks, customers, and partner organizations. Fraud detection systems addressing these issues with large-scale and field-wide transaction records are extremely helpful. However, developing effective and practical fraud-detection systems poses three major challenges. First, defining a learning target based on business or field requirements is crucial for achieving a truly useful system. Second, support from business expertise is essential for understanding features that are difficult to extract automatically, particularly those concerning certain fraud types. Finally, high-quality and genuinely large training sets are difficult to obtain.

The employment of Deep Learning as the main tool in fraud detection is particularly suitable for three reasons: (1) the availability of a truly large-scale transaction dataset, (2) wide enough temporal coverage of the dataset to capture both stable and evolving user behavior, and (3) the repetition characteristic of fraud types that permits formal modeling of fraud scenarios. A supervised-learning approach has been used to detect common types of fraudulent transactions at individual point-in-time locations. A recurrent deep-learning model has been designed to identify more complex fraud types based on human-usage characteristics over time. Furthermore, transaction-level features, engineered on the underlying transaction data with the support of domain expertise, have been employed to help in detecting various types of fraud including text-related and risk-related frauds.

Keywords—Deep Learning, Fraud Detection, Digital Banking, Transaction Fraud, Financial Cybersecurity, Anomaly Detection, Neural Networks, Transaction Monitoring, Financial Crime Detection, Intelligent Banking Systems, Artificial intelligence, Data engineering, Deep learning, Fraud detection, Recurrent neural networks, Supervised classification, Transaction data.

1. INTRODUCTION

With the growth of e-banking services, banks are exposed to fraud attacks such as forgery, account takeover attacks, money laundering, and hacking. Fraud detection in digital banking transactions has become a crucial area of research. Important precautions are necessary to protect banks and customers from irreversible financial losses because fraud detection systems in the real world have not attained high enough accuracy to assist financial fraud managers in responding quickly to transactions containing possible fraud patterns.

Intelligent fraud detection (IFD) systems use deep learning techniques and machine learning models built on transaction history pattern information to detect fraudulent transactions. A selected set of systematic review contributions shows that these methods provide promising strategies and yield excellent results in comparison with traditional systems. The systematic review also provides an overview of deep learning approaches used for intelligent fraud detection across different domains to detect such electronic fraud by classifying transactions as either fraudulent or legitimate.




Fig. 1. Publication years of the 30 cited references.

2. Overview of Digital Banking Fraud Landscape

Fraud in digital banking can emerge from various channels, including online banking, mobile banking, and social engineering scams, each hosting their own threats. Over its history, fraud has been mainly perpetuated by organized groups, but the arrival of cryptocurrencies has simplified the masking of the money laundering processes, drawing in a plethora of new actors. In this environment, anti-fraud systems must constantly adapt, utilizing and fusing the information gained from all available fraud signals. Fraud indicators lurk under the surface of legitimate transactions, making detection challenging. Moreover, detection systems are evolving, and the age of Machine Learning (ML) is set to revolutionize fraud detection.

Fraudulent transactions represent a small portion of the overall volume; hence, the greatest challenge in automatic detection lies in the imbalance of the available dataset. The volume of data, however, is such that, from a statistical perspective, those indicators can be detected and exploited for automatic systems. Digital banking fraud is a worldwide crime and a multi-channel threat for financial institutions. It generates significant losses at operating and financial levels. Online banking fraud in particular is a complex phenomenon: different threat actors exploit both the social trust in online banking services and the web channel's inherent vulnerabilities. Besides the financial losses, the impact on the fraud victims is considerable, and financial institutions strive to provide the highest level of protection for their customers.



Fig. 2. Overview of the supervised fraud-detection pipeline described in the paper.

3. Fundamentals of Deep Learning in Fraud Detection

Deep Learning fundamentally relies on technology that allows computers to solve complicated problems by building high-level abstractions from the data drawn on. Deep neural networks (DNNs) contain a massive number of parameters. Often, they are endowed with several hidden layers and possess great modeling powers, making them prone to overfitting. As a consequence, for a DNN to generalize effectively, it requires a huge amount of data. In addition to being data-hungry, DNNs are also recognized for their poor generalizability on unseen classes, that is, they are sensitive to class imbalance. Furthermore, the black-box nature of DNNs has raised serious challenges in their application to fraud detection in digital banking transactions. Despite the challenges, tremendous engineering efforts resulted in very accurate DNN implementations.

Nature of a deep architecture invariably determines the types of data that can be effectively processed. DNNs use a rich set of representations for naturally Class-structured data, while convolutional neural networks (CNNs) are recognized for the high-performance processing of image data. Recently, recurrent neural networks (RNNs), long short-term memory, and gated recurrent units have become prevalent in handling sequential data such as stock prices, text, and other time series. Indications also exist of the successful application of a temporal model to fraud detection, thereby detecting frauds in account-takeover and muleing. GCNs operate directly on graph-structured data yet. A GCN-based model detected muleing with good accuracy. Advantageously, the training time can be much shorter than for counterparts, which are not augmented with subgraphs during training.




Fig. 3. Leading publication venues among the cited references.

A. Supervised Learning for Transaction Classification

Supervised learning can be applied to binary or multi-class fraud classification. Supervising the classification process requires a fraud labeling scheme that defines the possible categories of a transaction (transaction-level labels), or the fraud class of a user at a particular time period (user-level labels). Supervision is generally achieved through the minimization of a cross-entropy loss, which measures the divergence between predicted probabilities and true target classes. Evaluation is commonly performed by calculating accuracy, precision, and recall on a separate validation hold-out set. The performance of supervised classification methods relies on the selection of relevant input features, which capture the transaction's characteristics and transactional patterns.

Training data for classification can stem from an attack vector considered as a potential point of failure. For instance, account takeover attacks can be fruitfully modeled through user-level labels, that discriminate between users that have been control of the how malicious actors FA system can be exploited and previous indicators of fraudulent activity on the same accounts.

Mathematical Formulas:

Cross-entropy loss for supervised classification:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [y_i \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)] \quad (1)$$

Evaluation metrics:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}, \quad F_1 = \frac{2 \text{ Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (2)$$

Amount relative to rolling statistics, and amount ratios:

$$z_t = \frac{a_t - \mu_{t,n}}{\sigma_{t,n}}, \quad \mu_{t,n} = \frac{1}{n} \sum_{k=1}^n a_{t-k}, \quad r_t^{(n)} = \frac{a_t}{\mu_{t,n}}, \quad r_t^{(1)} = \frac{a_t}{a_{t-1}} \quad (3)$$

Velocity feature over window Δ :

$$v_t(\Delta) = \frac{|\{j: t - \Delta \leq t_j \leq t\}|}{\Delta} \quad (4)$$

Great-circle (haversine) distance between user and merchant ($R = 6371$ km):

$$d = 2R \arcsin \sqrt{\sin^2 \frac{\phi_2 - \phi_1}{2} + \cos \phi_1 \cos \phi_2 \sin^2 \frac{\lambda_2 - \lambda_1}{2}} \quad (5)$$

Recurrent hidden state over a customer's transaction sequence:

$$h_t = \tanh(W_x x_t + W_h h_{t-1} + b), \quad \hat{y}_t = \sigma(W_o h_t + b_o) \quad (6)$$

Attention over the hidden states:

$$\alpha_t = \frac{\exp(e_t)}{\sum_k \exp(e_k)}, \quad c = \sum_t \alpha_t h_t \quad (7)$$

4. Architectures and Models

Architectures supporting fraud detection share commonalities with applications in other domains. Varieties defined by commercial architecture and transformation data sequence (e.g., image, audio, graph, textual) offer advantages and tradeoffs influencing fraud detection performance, scalability, interpretability, and defenses against adversarial behavior. Key architecture-related parameters include explainability (e.g., feature importance for end-users and regulators), temporal factors

(including freshness of training data), involvement of latent representation (whose role and reliability), model capacity and its scalability with incremental model evolution.

Success in fraud detection with neural networks suggests analyzing architectures explicitly designed for the task. Temporal signaling is frequently accomplished with recurrent layers capable of modeling long dependencies in classification and generation settings. Recurrent neural networks (RNNs) constitute the most generic sequential model. Models addressing gradient-vanishing/explosion include long-short-term memory (LSTM) networks, gated recurrent units (GRUs), and multilayer perceptrons shared among timesteps. Temporal convolution enables the exploitation of convolutions with sensitivity kernels wrapping in temporal direction. Modern attention models, effector modules physical-intuition-free and transformer architectures have dissimilar yet-interesting temporal behavior. An attention window facilitates feature selection in spatial-fading manner, in other settings training attention activations preserves full sequence conditioning. Applied to transaction sequence originated by a given bank account, RNNs and LSTMs/GRUs bridge semantic/finality gap, and attention model preserve short-distance semantic.

TABLE I. DEEP LEARNING MODELS DISCUSSED FOR FRAUD DETECTION

Model	Suited data	Role in fraud detection
DNN	Structured / class-labeled	Supervised transaction classification; data-hungry, prone to overfitting, black-box
CNN	Image-like / temporal kernels	High-performance pattern extraction; temporal convolution over sequences
RNN / LSTM / GRU	Sequential (time series)	Account takeover and muling via hidden state over customer history; mitigates vanishing gradient
Attention / Transformer	Sequences	Feature selection over a window; preserves short-distance semantics
GCN	Graph-structured	Mule-account detection on transaction graphs

A. Recurrent and Temporal Models for Sequential Transactions

Sequence modeling capabilities are useful for transaction classification since sequences of transactions are available for each customer, and the nature of transaction behavior is sequential. Customers perform transactions at different points in time, and not specifying the order of the transactions implicitly assumes that the transactions within a specific time window are independent and identically distributed. Temporal embeddings can be introduced to provide time-aware representations of the transactions such that the model can code different effects of different transactions along the time axis. When the fraud detection model is applied to sequences that contain transactions performed at interior points that were not part of the training set, the hidden state encouraging generalization can reduce the likelihood of false negatives. It is also possible to transfer the knowledge of performing normal transactions through a hidden state. It can be used to control-based generation and mark the generation of suspected fraudulent conducting within a time window.

Hiding state neural networks can also provide slightly different sequential representations. The model can either take the set of transactions in the whole time horizon without adding proper modeling of the time stamps or randomize the order of transactions until a small synthesis time window and use a hidden state to keep proxy information of all previous transactions in the window. When the synthesis time is large enough, the first approach emulates a recurrent neural network training step, but with a larger view because windows from different parts of the sequence can overlap. The windowed design utilizes all transactions, although some are lost when it is near the early or late ends of the long transactional sequence.

TABLE II. TRANSACTION-LEVEL FEATURE GROUPS

Group	Definition	Fraud signal
Amount (5)	Amount, z-score vs. rolling mean/SD, ratio to previous n and previous 1, rolling mean	Abnormal spending pattern

Velocity	Transaction count in last n s/min/h/days ÷ period	Rapid bursts of activity
Geolocation	Great-circle distance, user to merchant (km)	Location mismatch
Device	Browser user-agent fingerprint	Behavioral deviation
Merchant	Historical counts of transactions, unique users/devices	Merchant-specific volume anomalies
Temporal	Time since last transaction; time of day	Unusual timing

5. Feature Engineering and Data Engineering Considerations

Relevant features at the transaction level can add a further dimension to modeling efforts by enabling decision trees or tree ensembles to identify volume anomalies. Constructing aggregates along several dimensions (most commonly time) enhances the ability of nonsequential models to capture volume patterns. We adopt a transaction-level framework, which calls for careful consideration of the data construction pipeline.

Individuals make purchases and funds are transferred via several different types of financial instruments and channels (credit card, debit card, e-check, P2P transfer, retail transfer), using an assortment of devices (smartphone, laptop, e-tablet). Moreover, the characteristics associated with transacting with different merchants differ (e.g., an airline offers ticket sales at most a few times a day, while retail shops operate at a vastly different scale). Subset-specific volume aggregate features, combined with transaction- and device-fingerprint-centric features, attempt to account for these differences.

Moreover, the majority of the transactions in a credit card dataset are in a supportive fashion labeled “No Fraud.” Easiest to classify and the majority class, it can become a non-predictive feature, acting as noise in the learning algorithm. To alleviate this issue, volume pattern aggregate features significantly improve classification outcomes. The aim is to reduce the level of noise and provide a least-squares-fitting-like behavior to subsequent modeling efforts so that the “No Fraud” samples are not completely dismissed or ignored by a decision-tree-based algorithm.

A. Transaction-Level Features

Risk of transaction fraud can potentially be reduced significantly by taking advantage of additional information related to the transaction at the transaction-level. Although transaction amount is typically the most important feature across most domains, the pattern of amounts can also act as a good predictor of the likelihood of a transaction being fraudulent. Velocity features have also been shown to have high predictive power in the fraud detection context. Other features that have also previously been shown to be effective are geolocational information of the transaction, device fingerprinting, merchant behavior, and temporal information, all of which can potentially be extracted at the transaction-level.

Five features are proposed based on the transaction amount: the transaction amount itself, the transaction amount relative to a rolling mean and standard deviation, the current transaction amount divided by the previous n transactions amount, the current transaction amount divided by the previous transaction amount, and the rolling mean of the last n transactions amounts. Velocity features are generated by looking at a count of the transactions in the last n seconds/minutes/hours/days and taking the ratio of that count to the respective period. Geolocation features use the latitude and longitude coordinates of the merchant and the user and use the Great Circle distance (in km) between the two locations. Device fingerprinting features using the user-agent of the browser help identify the type of device used, and can be further used to identify if there is a behavioral deviation from the merchant. Merchant behavior features use the historical count of transactions and the historical unique count of users/devices for the merchant in the last n seconds/minutes/hours/days. Finally, temporal anomaly features look at the time difference from the last transaction as well as the time of day of the transaction.

6. Conclusion

Concerns regarding fraud and security have heavily retarded the growth of the digital banking sector, leading many users to question the overall security of their information. Without fraud detection and prevention, digital banking transactions can be difficult for all involved. This has become a factor for declining online users in this sector. Mahendra Mohan, Kiran P, and Fathima S. presume that intelligent systems will not only help customer service but also fraud detection and fraud mitigation. Detecting and classifying general fraud in real-time is still quite challenging because of the continuously changing environment and sophisticated attacker models. Feature analysis is a significant challenge in real-time fraud detection.

Fraud detection using transactions has become more efficient with the use of deep learning. Features are generated at a single transaction level with the assistance of deep learning architectures. ML could also solve the credit scoring and insurance underwriting problems where neither fraud has occurred nor is an account closed, but rather data have to be classified as indicated above. Successful supervised learning with transaction data is even more complicated than it appears at first glance. A fundamental premise of fraud detection is that the ML can determine, given the characteristics of the transaction compared to a user-defined normal state variable of a transaction, whether the present transaction is normal or fraudulent. The full understanding of what is present in data is a different process from the actual classification of data for fraud detection.

REFERENCES

1. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, 100402. DOI
2. Mattaparthi, R. (2021). Unified Data Lineage and Quality Governance Framework for Multi-Source Sensor Streams in Heavy-Duty Powertrain Manufacturing. *Online Journal of Mechanical Engineering*, 1(1), 1-15.
3. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*, 1(1), 1-13.
4. Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. *Universal Journal of Computer Sciences and Communications*, 1(1), 1-21.
5. Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 8, 151. Springer
6. Carcillo, F., Le Borgne, Y. A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331. ScienceDirect
7. Dang, T. K., Tran, T. C., Tuan, L. M., & Tiep, M. V. (2021). Machine learning based on resampling approaches and deep reinforcement learning for credit card fraud detection systems. *Applied Sciences*, 11(21), 10004. MDPI
8. Domashova, J., & Kripak, E. (2021). Identification of non-typical international transactions on bank cards of individuals using machine learning methods. *Procedia Computer Science*, 190, 178–183. ResearchGate
9. Domashova, J., & Zabelina, O. (2021). Detection of fraudulent transactions using SAS Viya machine learning algorithms. *Procedia Computer Science*, 190, 204–209. Ouci
10. Mukesh, A., & Aitha, A. R. (2021). Insurance Risk Assessment Using Predictive Modeling Techniques. *International Journal of Emerging Research in Engineering and Technology*, 2(4), 68-79.
11. Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479, 448–455. ScienceDirect
12. Forough, J., & Momtazi, S. (2021). Ensemble of deep sequential models for credit card fraud detection. *Applied Soft Computing*, 99, 106883. ScienceDirect
13. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. *World Journal of Clinical Medicine Research*, 1(1), 1-18.
14. Ileberi, E., Sun, Y., & Wang, Z. (2021). Performance evaluation of machine learning methods for credit card fraud detection using SMOTE and AdaBoost. *IEEE Access*, 9, 165286–165294. IEEE Xplore
15. Kim, E., Lee, J., Shin, H., Yang, H., Cho, S., Nam, S. K., Song, Y., Yoon, J. A., & Kim, J. I. (2019). Champion-challenger analysis for credit card fraud detection: Hybrid ensemble and deep learning. *Expert Systems with Applications*, 128, 214–224. ScienceDirect
16. Reddy, V. A. R. (2021). Challenges in Standardizing Member Eligibility Data Across Multi-Payer Healthcare Ecosystems. *International Journal of Medical Toxicology and Legal Medicine*, 24(3), 1-19.
17. Mekterović, I., Karan, M., Pintar, D., & Brkić, L. (2021). Credit card fraud detection in card-not-present transactions: Where to invest? *Applied Sciences*, 11(15), 6766. MDPI
18. Misra, S., Thakur, S., Ghosh, M., & Saha, S. K. (2020). An autoencoder based model for detecting fraudulent credit card transaction. *Procedia Computer Science*, 167, 254–262.
19. Mqadi, N. M., Naicker, N., & Adeliyi, T. (2021). Solving misclassification of the credit card imbalance problem using near miss. *Mathematical Problems in Engineering*, 2021, 7194728. ScienceDirect
20. Zhang, X., Han, Y., Xu, W., & Wang, Q. (2021). HOBA: A novel feature engineering methodology for credit card fraud detection with a deep learning architecture. *Information Sciences*, 557, 302–316. ScienceDirect

21. Loganathan, R. (2021). Integrated Risk and Compliance Frameworks for Global Data Center Operations: A Governance-Centric Approach. *Universal Journal of Computer Sciences and Communications*, 1(1), 1-26.
22. Carneiro, N., Figueira, G., & Costa, M. (2017). A data mining based system for credit-card fraud detection in e-tail. *Decision Support Systems*, 95, 91–101.
23. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784–3797.
24. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P. E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234–245.
25. Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2018). Credit card fraud detection using AdaBoost and majority voting. *IEEE Access*, 6, 14277–14284.
26. Roy, A., Sun, J., Mahoney, R., Alonzi, L., Adams, S., & Beling, P. A. (2018). Deep learning detecting fraud in credit card transactions. In *2018 Systems and Information Engineering Design Symposium (SIEDS)* (pp. 129–134). IEEE.
27. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113.
28. Bahnsen, A. C., Aouada, D., Stojanovic, A., & Ottersten, B. (2016). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications*, 51, 134–142.
29. Van Vlasselaer, V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems*, 75, 38–48.
30. Whitrow, C., Hand, D. J., Juszczak, P., Weston, D., & Adams, N. M. (2009). Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery*, 18(1), 30–55.
31. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602–613.