Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-416-0_7

Chapter 7: Fraud Detection Systems Using Deep Learning and Behavioral Intelligence

7.1. Introduction

Fraud Detection Systems Using Deep Learning and Behavioral Intelligence: The rapid adoption of e-commerce technology has produced unprecedented improvements in lifestyle and convenience of purchasing. Over the years the mechanisms for shopping have also been improved and automated for the customers' benefit. In the backend however, contending for such highly streamlined technological improvement are threats and vulnerabilities attacking the financial pillars of such systems. Regular security production systems handle detection based on rules or heuristics defined by the historical behaviour, these systems have very little or no impact against a new strike. Moreover these fraud production systems handling digital transactions typically operate based on known malicious signatures or detect distinctive patterns. But the unparalleled rapid shift in banking especially the digital economy has outgrown the surge of fraud detection. As new-age fraudsters employ more sophisticated approaches to attack, companies need to adapt new digital fraud detection techniques to counter these imminent frauds. Both enterprises and academia have started exploring detection methods underpinned by concepts like behavioural intelligence and deep learning with a major focus on supervised learning techniques. Such techniques can indeed enhance detection efficiency, nevertheless when done in isolation they fall short in addressing the critical imbalance and the data privacy issues. Hence a novel crime prevention system taking a combined approach with only supervised, semi-supervised or unsupervised learning invariably the core challenge of model drift during the operationalisation stage invariably determines the success or failure of the fraud detection. The performance metrics used to evaluate the proposed system predict false positives with grave importance since fraudsters usually stage a fraudulent attempt when they know for sure that the system is less likely to detect their attempt. With increased focus on the balancing ,trait-level feature engineering, real-time inference, model drift, monitoring and scalability gives directions for future research in fraud detection systems.

7.1.1. Overview of the Study

Advanced fraud detection systems employ supervised and semi-supervised deep learning models, anomaly detection, and unsupervised deep architectures. User and Entity Behavior Analytics detects fraud using behavioral intelligence, centered on user actions and transaction content. Feature engineering distills behavioral signals for supervised models. Data quality, privacy, and imbalance considerations affect reliability. Authorship verification, active learning, and hyper-parameter optimization influence performance. Real-time inference, scalability, model drift, retraining, drift detection, and governance are key operational challenges.

Fraud detection examines attributes of legitimate and fraudulent events for accurate categorization, monitoring, prevention, and security of targeted systems. Fraud is defined as a deceitful act thriving on imbalance, disorder, and determination, offering reward but no effort. Fraudulent activities appear similar to legitimate events, exploiting the very values that make detection difficult. Fraud evolves, taking different forms, approaching detection systems, and making detection ineffective and less workable. As novel consumer products and services appear, fraud Information Technologies increase, requiring greater Government policy and awareness. Therefore, Fraud Detection Systems are significant, discovering new forms of deception based on behavioral intelligence and deep learning.

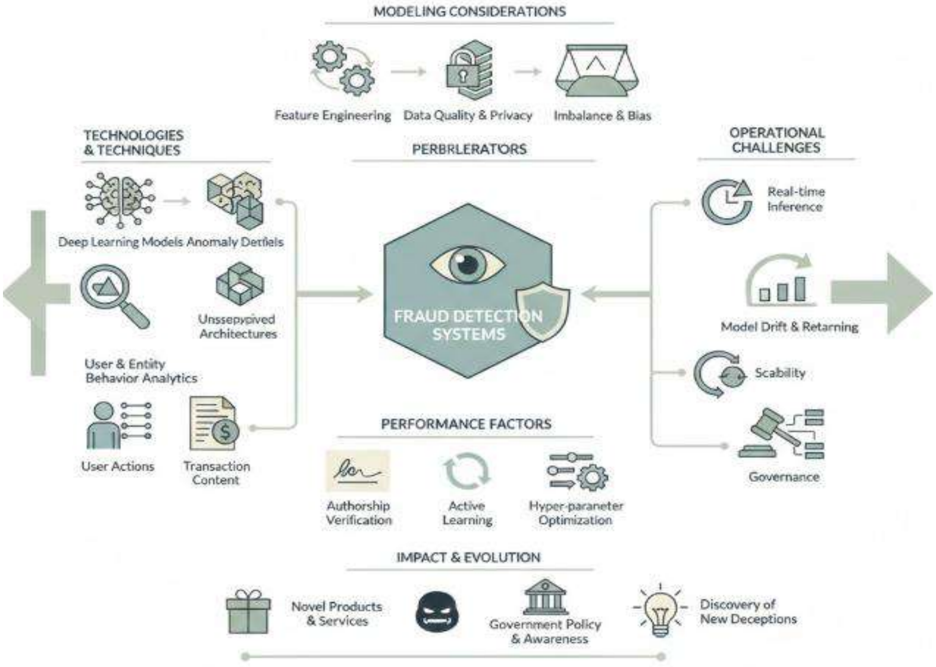

Fig 7.1: Dynamic Deception and Behavioral Intelligence: Leveraging Hybrid Deep Learning for Real-Time Fraud Detection and Governance

7.2. Foundations of Fraud Detection

Practical methods from graph theory, statistics, and machine learning are examined in terms of how they help to establish indications—a key part of the collaborative approach. Expectations with respect to fraud detection cannot be based solely on quantifying the attributes associated with fraud; such attributes can also be addressed by techniques such as data mining. Detection builds on a different cornerstone, one linking prediction to decision-making and behavioural analysis. The work discusses the key principle underlying fraud-detection techniques—the point of collation that determines the quality of the indication that a technique generates. Practical implementations of this principle are found in analytic techniques, such as anomaly detection and link-prediction approaches, as well as supervised models trained on data from previous frauds. A final group of operable models is composed of semi-supervised techniques collectively labelled user and entity behaviour analytics, which synthesise the aforementioned methods in a comprehensive governable deployment scheme.

Fraud detection requires an amalgamation of implementation decisions tailored to the specific problem domain. Any fraud-detection method usually involves rules, heuristics, or rational behaviour associated with appearance or propensity to commit fraud; these are not applied for detection but explicitly control processes such as alert prioritisation and indicator management. Indication is the degree of confidence that a fraud or attempt at fraud has occurred. Various techniques provide different degrees of indication depending on the quality of the data, the parameterisation of the technique, and the closeness of the detected behaviour to that of known fraud.

7.2.1. Key Principles of Fraud Detection Techniques

Fraud detection is a broad area of research in the field of information security. There are different types of fraud activities, including organizational fraud, account takeover, credit card fraud, money laundering, and malware distribution. Fraud detection systems filter out fraudulent activities from normal ones either by monitoring behavioral changes or by leveraging existing knowledge of the domain. The complexity of fraud detection is due to the desire of malicious users to deceive such systems and their skillful use of a variety of counter-speeches. In particular, the counter-speeches used in the fraud activity tend to be similar to the normal activities, which poses a great challenge to researchers to design an effective fraud detection system. Fraud detection has been addressed as a classification problem with labeled data and as an anomaly detection problem with unlabeled data.

Fraud detection can be classified as a supervised classification problem or an anomaly detection problem. In a supervised classification problem, the training data is represented using a set of features along with their corresponding labels. During the training phase, the model is learned based on the labeled training data, and during the testing phase, the model is tested using the unseen testing data. If the proportion of data points belonging to a specific class is very small in comparison to the rest of the classes in the data set, then the problem can be viewed as an imbalanced classification task. In an anomaly detection problem, the training data is represented using a set of features, and only a very small number of points belong to a specific class in the training data. The important focus of the anomaly detection algorithm is to detect the unseen suspicious points in the testing phase.

7.3. Deep Learning Architectures for Fraud Detection

Fraud detection is predominantly an iterative process in which domain experts design, implement, and refine algorithms that classify transactions as legitimate or fraudulent. Increasingly, these algorithms closely resemble traditional machine learning algorithms in which features extracted from transactions are fed into classifiers for label prediction. In these techniques, deep learning is a black-box approach in which domain knowledge can guide the preprocessing and engineering of features but plays a lesser role in the design of classifiers themselves.

Deep learning techniques for fraud detection can be broadly categorized into three classes based on the types of models deployed during fraud detection. The first class includes supervised models with full labels, typically a small percentage of frauds among vast amounts of transactions. The second class includes semi-supervised models combining labeled legitimate signals with unlabeled transactions. The remaining transactions constitute the vast majority of data and may thus help improve detection efficacy through anomaly classification techniques. The final third class incorporates no labels during training but exploits patterns of legitimate behavior to detect fraudulent activity.

7.3.1. Supervised and Semi-Supervised Models

Various supervised and semi-supervised deep learning architectures have been proposed for a range of fraud detection applications. Recent work focuses on financial transaction fraud, specifically for credit-card transactions in which real-time processing and low latency are fundamental. Credit-card fraud detection has been addressed as both a binary and multiclass classification task, considering a small number of synthetic-Normal (class 0) and normal (class 1) tokens while the majority of the data belong to the fraud-oriented class (class 2). Several deep learning-based methods were evaluated for countering three of the existing fraud attack strategies, namely, application fraud, account takeover, and card-not-present attacks. A sophisticated ensemble model combining a random oversampling technique to address class imbalance and the CNN–LSTM hybrid architecture outperformed competitors. Other studies apply Bayesian networks and radial-basis-function networks to the same credit-card transaction setting.

In the context of social network fraud, Nuisance Detection in Model-Free User Representations deals with the task of detecting malicious behavior in a social network in real-time, especially for spam activity. The proposed approach leverages the recent Domain Adaptation (DA) literature and expands the idea of domain adaptation to Nuisance Adaptation (NA): the goal is to specialize in classifying a specific class while treating the others as a nuisance. Recent work also exploited a semi-supervised Convolutional Neural Network–Deep Convolutional Neural Network (CNN-DCNN) model and a deep hierarchical network with parallel Naïve Bayes layer for social spam detection.

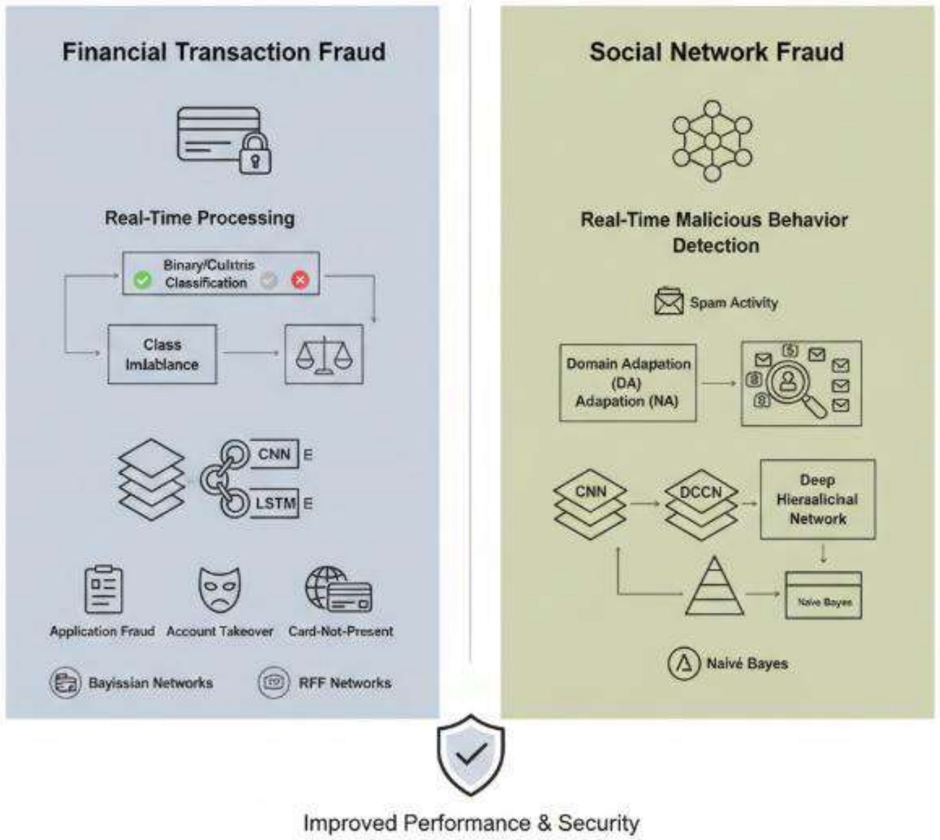

Fig 7.2: Multi-Domain Fraud Mitigation via Hybrid Deep Learning: From CNN–LSTM Financial Transaction Surveillance to Nuisance Adaptation in Social Networks

7.3.2. Anomaly Detection and Unsupervised Techniques

Anomaly detection techniques in fraud detection differ from traditional supervised learning methods. The main challenge is that labeled data (i.e., fraudulent samples) is rare compared to unlabeled data. In anomaly detection approaches, models are trained using normal data (i.e., non-fraudulent samples), and deviations from this normal behavior are considered fraudulent. In unsupervised learning, models are trained without labels. Attention-based autoencoders are now being used as a novel approach. Other unsupervised techniques—support vector data description, isolation forest, one-class support vector machine—use distances in the feature space to indicate whether a sample is similar to or different from the majority of samples. These techniques treat each sample in the data set as an anomaly detector by measuring distances and establishing an acceptance region with no samples.

While the majority of fraud cases may be labeled as rare events, Goutte and Gaussier proposed a new evaluation method with the principle that false negative errors are always more serious than false positive errors. This approach works even for strongly unbalanced data sets. In contrast to most traditional learning approaches that indicate a very low precision, a new representation measures the quality of a learning method with respect to its applicability in a real-world fraud-detection problem where undesirable deposits of collateral information represent a major problem. Other approaches include semi-supervised feature selection and novelty detection through Stone's separation theorem-based kernel methods for a random subset of samples.

7.4. Behavioral Intelligence in Fraud Prevention

Fraud detection technologies have incorporated user behavioral models reliably for years. User and entity behavior analytics (UEBA) adds another layer of analysis by using machine learning to establish the usual patterns of system entity behavior and detect deviations from these patterns. These deviations are critical because they often lie beyond the detection capabilities of traditional security controls and can thus indicate furtive adversarial activity. UEBA seeks to model user behavior for systems that typically involve many users, such as information-sensitive environments, email systems, VPNs, and cloud services. Similar to UEBA, user behavior analytics analyzes individual user accounts rather than many at once. This technique identifies unauthorized access to user accounts, usually through the collection of logs associated with the user and those on the user's account.

Despite these successful applications, the signatures generated still overfit the data. Greater understanding of binding-based craving will support the shaping of behavioral, trait, and script features that are more general to the domain. In the meantime, an alternative means of generating behavioral signals lies in the construction of features that represent an aggregation of user actions over different temporal windows. Behavioral signals built in this manner provide a preliminary assessment of whether a subject shares similar taste and purchasing patterns with peers. Like product affinity signals, behavioral signals are typically noisy because it is also common for users in any dataset to differ vastly in taste at some local level. However, when pooled over large enough periods, these signals are better captured. The prime candidates for feature engineering revolve around common behavioral analytics measures.

7.4.1. User and Entity Behavior Analytics

Behavioral intelligence, the use of data describing how people and systems typically behave, is integral to fraud prevention. User and Entity Behavior Analytics (UEBA) focuses on detecting threats by establishing baselines of behavior for specific entities. Anomalous patterns, generated through machine learning on many disparate signals, form a digital signature for future behavioral activity. Intrusion Detection Systems have long applied heuristics to identify when incoming activity does not conform to expected patterns. UEBA extends this thinking to a wider observation space.

UEBA grows in relevance as attacks increase in sophistication. Signature-based detection schemes by definition rely on earlier successful detection of an attack pattern. While risks of malware, phishing, and insider threats remain ever present, the likelihood of facility by clever attackers becomes higher as the number of known techniques expands. Increasingly destructive attacks, like those involving ransomware or distributed denial of service, especially call for further protection. The rising incidence of breaches as a result of credential theft and lateral movement techniques similarly reinforces the need for monitoring outside the traditional controls such as firewalls or anti-virus software. Marty et al. (2014) observe that existing system access controls normally delineate the boundaries of security, whereas UEBA focuses attention on what is happening inside those controls by examining how entities within the perimeter typically behave.

7.4.2. Feature Engineering for Behavioral Signals

Incorporating features that indicate behavioral changes or anomalies relative to history is key in fraud detection. Such information can be expressed in labels for a supervised model or, ideally, mapped into a feature space in conjunction with data from other users containing comparable behavior (or both). For instance, in malware detection, atypical patterns of API calls can be tracked using the temporal difference of a combination with the denial probability (assessed based on behavior of all other entities), thus alleviating the impossible-task aspect of searching for domain knowledge on malwares. In User and Entity Behavior Analytics (UEBA), some of the group-agnostic features might include session time (too short or long compared to previous sessions), distance between logins, logins from different countries (spatial mobility), and/or access to sensitive information or/or with unusual frequency.

In financial fraud detection, similar group-agnostic features can capture behavioral changes related with the current transaction: behavior of transactions from the same user, of transaction of the same amount/merchant, of transaction in a short time, and of transaction in a long period of time. An advanced tabular data representation allowing for feature sharing among users and accounts having correlation in terms of fraud and natural transactions is currently getting attention from the financial community. Finally, good performance can be obtained by train a model with both the original feature space (e.g., all transactional data, ATM or POS data) and behavioral feature.

7.5. Data Considerations and Evaluation Metrics

Fraud Detection Systems Using Deep Learning and Behavioral Intelligence: Data Considerations and Evaluation Metrics

The effective use of data is crucial for the success of any fraud analytics model. While it is relatively simple to capture the necessary data and train sophisticated deep neural networks, the power of a model lies not merely in its architecture but, more importantly, in the way that it uses the available data. For fraud detection tasks, consideration of data-related aspects is even more important, with specific emphasis on data imbalance and data privacy in addition to putative user behavior data. Moreover, the right evaluation metrics must be used to assess a model's performance.

Data quality, data balance, and data privacy are fundamental considerations for any machine learning model. Most fraud-use-case data have class imbalance, with a negligible fraction of fraud records in datasets containing a huge amount of non-fraud activities, and simple usage of accuracy cannot yield helpful insights about the performance of classifiers for detecting fraud. Considering the balance of data, precision and recall, precision-recall curve, AUC-ROC, and F1 score are more appropriate evaluation metrics for fraud detection analytics. In recent years, additional evaluation metrics considering the financial impact of fraud detection systems have also come into focus, assessing the efficiency of fraud detection in balancing cost versus risk.

7.5.1. Data Quality, Imbalance, and Privacy

Fraud Detection Systems Using Deep Learning and Behavioral Intelligence

Apart from the fundamental characteristics of the problem and the chosen approach, the quality and distribution of data have a critical impact in the experimentation phase, mainly due to common issues related to data quality, class imbalance, and privacy, especially in the context of real-world applications.

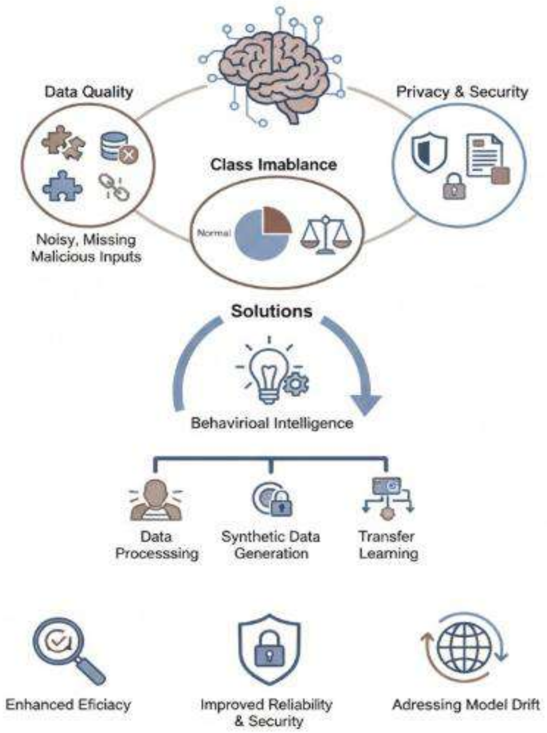

Fig 7.3: Navigating Adversarial Landscapes in Fraud Detection: Addressing Extreme Class Imbalance, Data Toxicity, and Regulatory Privacy Constraints

Data quality must be appropriate for the problem at hand. In fraud detection, when data come from the Internet, erroneous and malicious inputs are common. In this context, the gaps may not be due only to technical reasons, but also by design, since Internet data are used for malicious purposes (e.g. web pages spreading false news and private messages being misused for the production of mass calamity). Noisy, missing and/or irrelevant information negatively impact any machine learning algorithm, be it supervised or unsupervised. Furthermore, the scarcity of quality labeled data in fraud applications negatively impacts the efficacy and reliability of learning, validation, and testing processes.

Data class imbalance is a hallmark of fraud detection problems. It is rare to find formally labeled and validated benchmark datasets, since the institutions that deal with these problems protect their data to ensure the security of their customers and avoid further crime. In supervised learning, the normal class is heavily overrepresented, while the fraudulent class is very small and, in some scenarios, even much smaller than in a typical semi-supervised learning setting. In addition, the transference of the trained models also deserves careful consideration, as the distributions of the data may be different, leading to model drift. In fraud detection problems, privacy is a main concern due to legal restrictions imposed by the General Data Protection Regulation (GDPR) in Europe and similar laws in other countries.

7.5.2. Evaluation Metrics and Validation

Data quality, imbalance, and privacy—or a lack thereof—are some of the most important factors to be taken into account when developing fraud detection systems. The choice of an appropriate evaluation metric, as well as validation of the system, must also be emphasised, bearing in mind the consequences of false positives and false negatives and the system's target application.

Fraud detection is an imbalanced classification problem, in which the normal cases far outnumber the fraudulent ones. Therefore it is important for a metric to take into consideration the distribution of the classes in order to avoid being biased toward the majority class. Since the cost of a false positive is typically higher than that of a false negative in fraud detection, careful attention must also be paid to the selection of a suitable metric. Validating performance objectively and correctly is critical, given that any fraud detection system will ultimately be a classifier at its core and should be treated accordingly. A set of hold-out one-off datasets, which mimic the behaviour of the underlying data over the duration of the entire supervised learning or unsupervised anomaly detection process, is therefore recommended to avoid overfitting and selection bias.

7.6. Deployment, Monitoring, and Operational Challenges

The ultimate goal of fraud detection systems is to identify fraudulent activity as it occurs. Therefore, real-time inference and scalability are crucial. Further, because these systems learn behavioral patterns, data drift can occur even under normal business operations, necessitating regular feature retraining, selection, and derivation, along with continuous monitoring and governance of the entire pipeline.

Fraud detection is particularly challenging to scale because the financial costs and reputational damages incurred by any given organization are often not born by the organization itself, but by banks and insurers. The expected loss associated with each fraudulent activity is very low when compared to the scale of legitimate transactions. In response, organizations typically focus their operational resources on detecting and mitigating fraud scenarios that yield the highest expected loss, leaving other scenarios that are detected less frequently without sufficient operational resolution. Such an operational environment is ideal for the use of deep learning.

7.6.1. Real-Time Inference and Scalability

Fraud detection systems are increasingly expected to support real-time inference for a wide array of use cases. As transaction volumes continue to grow, query rates for dedicated fraud detection systems can reach hundreds of thousands of calls per second, and the low latency required for most applications directly translates into strict limits on the time-consuming operations for inference. The logical architecture of the fraud detection pipeline, including the interplay between batch and real-time scoring, has a strong impact on the computational load and response time of the models. A microservices architecture makes it possible to distribute the prediction load across multiple independently deployable services and improve response time, that is, the time it takes to produce a response after the query is received, by distributing heavy operations in the batch mode, only executed once or a few times a day.

Several anti-fraud use cases involve historical data all related to the same entity (e.g., a user, a credit card, etc.) and should therefore provide the user with a coherent 3D experience. For example, monitoring the risk of credit card stolen-purchase action requires information about past purchases made by a specific card on its online shopping profile, together with the other cards held by the same user, and also purchases sponsored by this user in the stores monitored. Fraud detection models should therefore also support multi-query scoring. From a functional perspective, it means that all the predictions related to the same entity should be processed together, allowing the construction of a single output containing all the results.

7.6.2. Model Drift, Retraining, and Governance

The practical application of a fraud detection model involves additional challenges beyond training. A profound understanding of the operating conditions helps to define requirements for real-time performance and the model's evolution over time. In real-world data, the same entities are observed repeatedly, which may result in predictable changes for a fraudster, and therefore, the model could stop being effective at some point. Control of model drift and operational signals from performance and business impact may therefore indicate the need for retraining. Ideally, operational control encompasses policies and procedures to ensure that all these aspects are assessed and that the fraud detection system evolves.

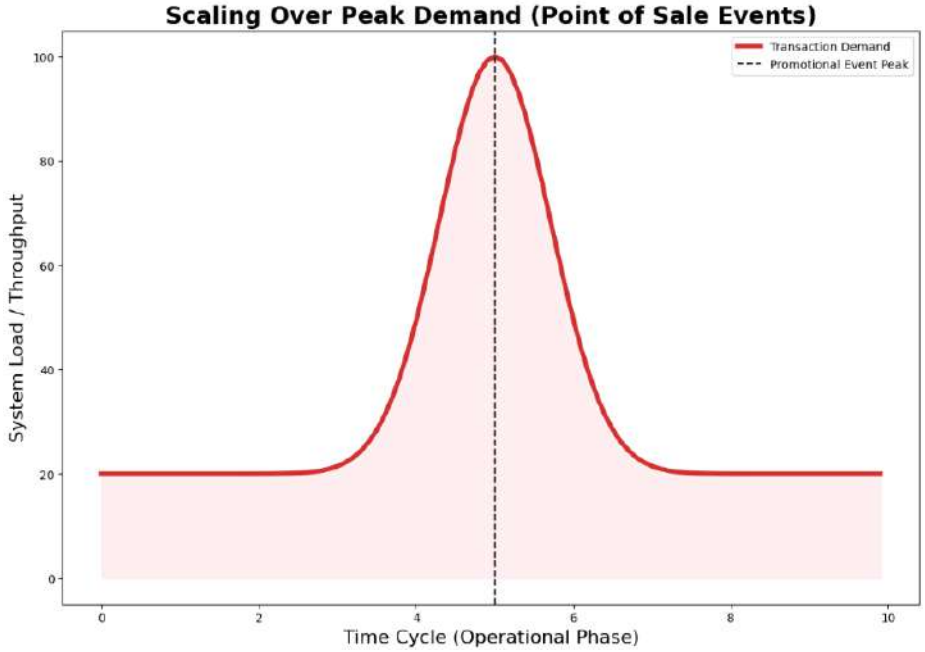

Fig 7.4: Scaling Over Peak Demand (Point of Sale Events)

Real-time inference of fraud detection models is a demanding task due to the requirements for low latency and high throughput. When the training and inference tasks are performed together in a single pipeline, the prediction latency should be under a few milliseconds, with throughput in the thousands to millions of predictions per second range. Furthermore, the operational deployment should scale over increases in demand, such as during points of sale of promotional products. Scalability usually relies on the implementation of distributed systems that can add new machines to the processing pool; in simpler implementations, horizontal scaling is often delegated to cloud providers.

7.7. Conclusion

This study demonstrates relevant aspects and advances of fraud detection systems using deep learning and behavioral intelligence. Despite the inherent difficulty of collecting and quality-assessing labelled fraud data since they are rare events, imbalanced and often need privacy protection, deep learning models applied to credit card transactions, insurance claims, cyber fraud, and other domains possess good potential to detect fraudulent activities. Evaluation results show that adding behavioural intelligence through User and Entity Behavior Analytics (UEBA) or behavior-aware feature engineering provides extra useful information. Operational aspects are equally discussed; real-time inference and scalability are confirmed as necessary for fraud-detection applications; model drift is a main concern under operational settings and retraining strategies are a common practice; governance is critical since the consequences of false positives and false negatives are experienced at many different business levels.

Research interest and publications on fraud detection are steadily growing. Attention has naturally shifted towards advanced approaches with potential for improving the detection quality, and selected research articles in fraud detection are hereby reviewed and discussed. The detected trends and recent works using deep-learning architectures for fraud detection incorporate concepts of behavioral intelligence by means of User and Entity Behaviour Analytics (UEBA) or misuses actual labelled data by exploiting behaviour-aware attributes derived from auxiliary problems. Suggestions for future work are given, with the evident need of other aspects in fraud-detection systems.

7.7.1. Final Thoughts and Future Directions in Fraud Detection

The evolution of the discipline has been impressive due, among others, to the massive amount of labeled data made available by catastrophic industrial breaches, the fast-growing computational power, and the capability to advance the state-of-the-art through gradual incremental improvements in the architectures and training paradigms. Nonetheless, even with success stories and several models in production, current efforts are highly fragmented and uneven.

In terms of volume, deep learning for fraud detection remains dwarfed by the many papers utilizing deep learning for vision and natural language processing tasks. Furthermore, the most deployed fraud detection models rely on classical machine learning instead of deep learning-biased strategies. Two main factors have hindered the massive adoption of deep learning in production environments. First, the scarce volume of labeled data in most organizations still makes anti-fraud models a challenging supervised problem rather than a well-posed semi-supervised or purely unsupervised task, disfavoring architectures delving in entirely end-to-end semi-supervised and unsupervised training paradigms. Second, the speed, scalability, and interpretability constraints of deployed production models for real-time user decisions are ad-hoc solved by specific solutions tailored for classical machine learning models, hence by-passing the usage of more sophisticated neural network architectures.

Behavioral intelligence helps stakeholders achieve better security by incorporating user and entity behavioral analytics in their fraud prediction systems. On the one hand, the human-centered nature of fraud leads to the analysis of genuine user behavior as a means of supporting decision-making. On the other hand, understanding users and entities over time is practiced and maintained by humans, which ought to be reflected in fraud prediction techniques. Considering both aspects, developing user and entity behavior models not only generates a plethora of behavioral information to be fed into anti-fraud models, but offers data streams capable of triggering alerts for suspicious activities.

References

Dal Pozzolo, A., Bontempi, G., Snoeck, M., & Snoeck, M. (2020). Adversarial drift detection in credit card transactions. Pattern Recognition Letters, 136, 252–258.

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.

Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 653-674.

Nagabhyru, K. C. (2023). From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust. Available at SSRN 5697663.

Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. Expert Systems with Applications, 240, 122156.

Kaggle, M., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Calibrating probability with undersampling for unbalanced classification. Pattern Recognition Letters, 140, 224–231.

Guntupalli, R. (2025, June). AI-Powered Data Analytics in Cloud Computing. In International Conference on Data Analytics & Management (pp. 280-289). Cham: Springer Nature Switzerland.

Zhang, Z., & Zhou, Z.-H. (2021). Adversarial learning for generative models: A survey. IEEE Transactions on Knowledge and Data Engineering, 33(3), 1026–1040.

Varri, D. B. S. (2022). AI-Driven Risk Assessment And Compliance Automation In Multi-Cloud Environments. Available at SSRN 5774924.

Carcillo, F., Dal Pozzolo, A., Bontempi, G., & Snoeck, M. (2021). Scarff: A scalable framework for streaming credit card fraud detection with concept drift adaptation. Information Fusion, 71, 182–197.

Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. International Journal of Scientific Research and Modern Technology, 1(12), 227–237. https://doi.org/10.38124/ijsrmt.v1i12.1111

Juszczak, P., Adams, N. M., Hand, D. J., Whitrow, C., & Weston, D. J. (2008). Off-the-peg and bespoke classifiers for fraud detection. Computational Statistics & Data Analysis, 52(9), 4521–4532.

Rao, A. N., Garapati, R. S., Suganya, R. T., Kaliappan, A., & Kamaleshwar, T. (2025, August). Smart Solar Harvesting and Power Management in IoT Nodes Through Deep Learning Models. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

Bahnsen, A. C., Aouada, D., & Ottersten, B. (2015). Cost-sensitive decision trees for fraud detection. Expert Systems with Applications, 42(3), 977–986.

Vadisetty, R., Polamarasetti, A., Goyal, M. K., Rongali, S. K., kumar Prajapati, S., & Butani, J. B. (2025, May). Cloud-Based Immersive Learning: The Role of Virtual Reality, Big Data, and Generative AI in Transformative Education Experiences. In 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC) (pp. 1-6). IEEE.

Wang, S., Cao, J., Yu, P. S., et al. (2022). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 54(2), 1–38.

Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.

Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. ACM Computing Surveys, 52(1), 1–38.

Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.

Raghavan, S., & Manchanda, P. (2021). Behavioral analytics for fraud detection. Journal of Marketing Analytics, 9(2), 73–89.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765–4774.

Siva Hemanth Kolla. (2022). Knowledge Retrieval Systems for Enterprise Service Environments. International Journal of Intelligent Systems and Applications in Engineering, 10(3s), 495–506. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/8037

Goodfellow, I., Bengio, Y., & Courville, A. (2020). Deep learning. MIT Press.

Zhou, C., Paffenroth, R. C., et al. (2017). Anomaly detection with robust deep autoencoders. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 665–674.

Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.

Chen, Y., Kou, G., Peng, Y., & Alsaadi, F. E. (2021). A multi-layer ensemble framework for credit risk prediction. Expert Systems with Applications, 174, 114762.

Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the workforce. Journal of Risk and Insurance, 89(2), 1–38.

Xu, Y., Sun, J., & Liu, J. (2023). Fairness-aware machine learning for insurance risk prediction. IEEE Access, 11, 24501–24513.

Inala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective.

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). Statistical and machine learning forecasting methods: Concerns and ways forward. PLOS ONE, 17(3), e0265480.

Richman, R., & Wüthrich, M. V. (2021). A neural network extension of the classical chain-ladder model. Insurance: Mathematics and Economics, 99, 331–347.

Serrano-Guerrero, J., Herrera-Viedma, E., & Olivas, J. A. (2021). Fuzzy and machine learning models in insurance risk classification. Applied Soft Computing, 102, 107070.