

RAG-Enhanced LLM Framework For Real-Time Fraud Detection In Financial Systems

Satishkumar Rajendran

University of Central Missouri, Warrensburg, USA

Abstract

Financial fraud remains a significant and evolving challenge and has become a multifaceted problem in modern financial systems, particularly due to the spread of digital payment systems and online financial services around the globe. Conventional rule-based fraud detection systems, which rely on fixed rules, are less effective in detecting complex and other rapidly changing patterns of fraud. The vast majority of the recent advancements in artificial intelligence have also seen the introduction of machine learning and deep learning algorithms that can analyse huge amounts of financial transactions and track abnormal behaviours. However, most existing fraud detection approaches rely primarily on structured transaction categories and cannot put into practice contextual data revolving around financial action. This review examines the rising prominence of large language model frameworks with retrieval as a part of real-time fraud detection in financial systems. The literature in the machine learning, graph analytics, and knowledge-enhanced language models' area is critically analysed to identify methodological trends, computational architectures, and performance characteristics of the new systems of fraud detection. As it may be observed with comparative analysis, ensemble machine learning models, deep neural networks and graph neural networks have significantly improved detection accuracy compared with classical statistical methods. Nonetheless, these techniques typically suffer a weakness when it comes to the ability to integrate unstructured data and contextual information pertinent to financial fraud investigation. Retrieval-augmented architectures can provide this answer since during the inference process, fraud detection systems view external sources of knowledge as dynamic repositories that can be consulted in the process. The latter systems integrate contextual information on financial behaviour with transaction analytics, and, furthermore, can enhance detection accuracy, interpretability, and adaptability to new schemes of fraud. The major research challenges include scalability in high-frequency transaction settings, combining heterogeneous data sources, explainability of model results, and privacy of sensitive financial data. Further research in the hybrid systems of graph analytics and retrieval-augmented reasoning systems and streaming analytics systems may help to develop the next generation fraud detection systems, which can meet the ever-growing financial threat.

Keywords: Retrieval-Augmented Generation (RAG); Financial fraud detection; Graph neural networks; Large language models; Transaction anomaly detection .

1. Introduction

Financial fraud has become one of the most persistent and technically sophisticated threats in the modern financial frameworks. The rapid digitization of financial services, growth of online payment systems, and adoption of decentralized technologies all increase the attack surface for malicious actors. Fraudulent activity has increased significantly over the past decade, causing banking institutions and payment systems to incur billions of dollars in annual losses from credit card fraud, identity theft, account takeovers, insurance fraud, and money laundering. Traditional detection systems have struggled to keep pace with rapidly evolving fraud patterns. Consequently, intelligent fraud detection frameworks have emerged as a key research issue in the financial technology, machine learning, and cybersecurity spaces [1].

Initial computational fraud detection methods were based mostly on statistical anomaly detection and rule-based decision systems. A common way of building these systems was

using hand-crafted rules based on the domain knowledge and historical trends of fraud. Although rule-based systems showed moderate performance in financial monitoring systems at the early stages, their weakness was witnessed as more financial transactions and fraud schemes emerged. Detection structures based on rules tend to have high false positives, are not flexible towards new types of fraud and demand a lot of manual maintenance. As a result, machine learning systems have begun to be used for automated fraud classification and anomaly detection of financial data sets [2].

The models of fraud detection based on machine learning brought the capability to process large volumes of transactional data and detect more complex nonlinear patterns that cannot be easily detected with fixed rules. Logistic regression, decision trees, support vector machines and ensemble modelling techniques have been broadly used to handle fraud detection. More recently, deep learning models, such as convolutional neural networks, recurrent neural networks, and graph neural networks, have shown significant improvement in detecting fraudulent transactions in highly imbalanced data [3]. Despite these developments, current models still face several limitations in real-time financial monitoring systems. Such limitations are challenges with integrating outside sources of knowledge, weak interpretability, and inability to keep up with constantly changing fraud tactics.

The large language models (LLMs) have become an innovative technology in the field of artificial intelligence, and it has been demonstrated to be useful in reasoning, contextual perception, and representation of knowledge. More recently, the application of LLMs has been mentioned beyond natural language processing activities and into other fields such as cybersecurity and financial analytics, and decision support systems. LLMs have shown promise in the analysis of unstructured financial reports, suspicious metadata transaction detection and automated compliance observance [4]. However, as far as the direct application of the LLMs to real-time fraud detectors is concerned, they face several limitations, including hallucination risk, latency constraints, and inability to access more up-to-date domain knowledge.

The concept of retrieval-augmented generation (RAG) has become a highly promising process that allows increasing the reliability and context sensitivity of large language models. The concept of RAG systems involves neural language models and external retrieval processes, wherein the model itself has the ability to retrieve the relevant information in a structured repository, document stores or in a knowledge graph of a domain dynamically. It is an architecture that facilitates response generation based on the retrieved evidence as opposed to basing it on internal model parameters only. RAG frameworks have the possibility to combine transaction data with historical fraud patterns, regulatory guidelines, and external sources of intelligence into a single analytic pipeline in the context of financial fraud detection [5].

The combination of RAG architectures with fraud detection systems offers several advantages to the real-time financial monitoring. First, the incorporation of ever-changing sources of knowledge due to retrieval mechanisms can enable detection systems to adapt to the new fraud tactics. Second, explainability could be enhanced by retrieval-based reasoning as it could be used to provide traceable evidence to support a fraud classification decision. Third, RAG systems have a modular architecture, allowing integration with the existing machine learning models, rule engines, as well as risk assessment pipelines. These capabilities make RAG-enhanced LLM models particularly attractive to financial organizations which will need plug-and-play and flexible fraud detection models [6].

Despite this promise, there are technical challenges in the implementation of RAG-based fraud detection systems. Financial transaction streams have extremely tight latency

requirements and often the decisions are required to be made within a few milliseconds. The retrieval and reasoning components in RAG architectures should therefore be maximized with the aim of achieving real-time processing. The properties of financial data also include the extreme class imbalance in which the transactions involving frauds tend to take less than 1 percent of the total transactions. Under these circumstances, detection systems based on RAG require models to be carefully designed and tested. Also, other issues with financial AI systems include privacy of data, conformity with regulations, and adversarial interference [7].

Recent studies have begun to investigate the hybrid models that can be used to integrate graph-based fraud detection models and retrieval-enhanced reasoning systems. The relational structure of financial transaction networks is complex and is modelled successfully using graph neural networks or heterogeneous information networks. These representations added to the RAG-based LLMs may provide heightened contextual decision-making information on transaction behaviour by bridging to larger financial intelligence bodies. Nevertheless, there is a dearth of empirical research that assesses such hybrid architectures and the scalability of these systems in a production environment has not been properly determined [8].

The other critical issue relates to how interpretable the decisions that the advanced AI models produce can be when it comes to detecting fraud. The financial institutions are regulated by supervisory authorities that need auditable decision-making procedures. Conventional deep learning systems tend to be black-box systems and therefore it is not easy to understand why certain transactions were identified as fraudulent. This weakness can be partially addressed by RAG-based models, as they can provide retrieval-based evidence to support model outputs. Nevertheless, minimal research has been done on systematic analysis of interpretability mechanisms within systems of fraud detection based on RAG [9].

As the area of artificial intelligence technologies actively develops, and the complexity of the fraud schemes is increasing, there is an obvious necessity to provide a comprehensive analysis of RAG-enhanced LLM schemes to identify frauds. The available literature is in different domains including financial analytics, machine learning, information retrieval, and cybersecurity. Still, the studies remain fragmented, and a consistent perspective of RAG-enabled fraud detection systems has not yet been completely and clearly built in the academic body of literature.

This review critically examines the emerging application of retrieval-augmented large language models as a component of real-time financial fraud detection systems. The discussion is based on both the methodological innovations, computational models, and the empirical findings that are presented in peer-reviewed papers. Specific focus is on the details of combining the systems of retrieval, machine learning models, and financial transaction analytics. The review shows the main technical problems, methodological shortcomings, and future research opportunities in this rapidly evolving field by using a systematic review of the existing research trends.

2. Literature Review

Various methodological steps have been taken towards detecting financial fraud as financial systems become highly digitalized and interrelated. The first computational methods were based on statistical anomaly detection models that were aimed to detect abnormal transactional behaviour in structured data sets. Gradually, the opportunities in machine learning, deep learning, graph analytics, and natural language processing have added to the number of analytical tools that may be used to detect fraudulent behaviour in financial networks. Recently, retrieval-augmented large language models have been

developed as one such promising paradigm of combining heterogeneous sources of financial knowledge with machine learning-based pipelines to detect fraud. The current section critically reviews these developments through research contributions of significance to the fields of statistical modelling, machine learning frameworks, graph-based systems of fraud detection, and novel retrieval-based frameworks.

The major areas of activity of early fraud detection research were based on statistical learning methods used on transactional data. The dominant approaches included logistic regression and Bayesian classification models because they are interpretable and efficient to calculate. Bolton and Hand also showed that unsupervised anomaly detection algorithms may be useful in detecting deviant financial activity, in terms of comparing observed transactions to learned probability distributions [10]. Such models did not, however, represent nonlinear relationships well in high dimensional financial data. Fraud patterns are also likely to evolve rather quickly and so reduce the usefulness of the applied models that were trained on the fixed historical data, in the long run.

In order to overcome these restrictions, scientists started using data mining and machine learning techniques that can learn nonlinear decision boundaries. Random forests and gradient boosting algorithms, which are techniques of ensemble learning, have become common in fraud detection systems. Whitrow et al. proved that aggregation methods of transactions along with random forest classification increased the performance of fraud detection in credit card data significantly [11]. Ensemble models are especially useful in detecting fraud since they compose a combination of various decision rules and minimize overfitting on very imbalanced datasets. However, even state-of-the-art machine learning models need to do a lot of feature engineering and usually do not take into account contextual financial information beyond structured attributes of the transactions.

The deep learning solutions have also increased the features of the fraud detection systems as they allow the automatic extraction of features using large volumes of financial data. LSTM networks and recurrent neural networks have been extensively used in order to model the sequence of transaction in time. Jurgovsky et al. showed that models based on LSTM were able to provide sequential dependencies when in a transaction stream, enhancing the accuracy of detection of fraudulent credit card use [12]. Convolutional neural networks were also modified to do fraud detection activities, where one considers transactional data as structured matrices. The methods enable finding the spatial correlations between transaction features to result in better classification results as opposed to conventional machine learning models.

Graph fraud detection models are becoming increasingly popular because financial transactions are inherently relational. Many financial frauds are generally organized in a way such that they cut across multiple accounts, devices, or organizations, and create highly elaborate webs of interactions. Graph neural networks (GNNs) have been especially useful in the models of such relationships. Wang et al. illustrated that the graph convolutional networks used on the transaction networks are capable of detecting suspicious groups of accounts by spreading the relational information among the nodes [13]. The graph-based models can model higher order dependencies and community structure, which are challenging to identify with a traditional feature based model.

The other significant development is heterogeneous information networks which combine various types of financial institutions, like accounts, merchants, devices and transaction channels. Such networks enable fraud detection models to identify the cross-domain relationship and identify coordinated campaigns of fraud. Dou et al. suggested a graph neural network model that was aimed at detecting camouflaged fraudsters in financial networks using relationship dependencies among entities [14]. The suggested

model was found to be more resilient to adversarial conduct in financial transaction networks.

Although graph-based models have proven successful, several challenges remain in their implementation in real-time fraud detection mechanisms. Graph neural networks are associated with massive numbers of computationally intensive message-passing operations in large networks, which may be a latency bottleneck in high-rate financial systems. Besides, the majority of the graph-based fraud detectors can rely only on structured transaction data, and cannot use unstructured sources of financial intelligence, such as regulatory reports, fraud case studies, or a description of suspicious transactions.

The new possibilities of the analysis of unstructured financial data with the use of large language models have become a reality. It has been demonstrated that transformer-based architecture contains good contextual reasoning and representation of knowledge. Financial analytics have also applied transformer models to financial sentiment analysis, the predictive modelling of risks, and document classification. Yang et al. demonstrated that the extraction of risk indicators using financial disclosures and regulation filings can be performed with language models that are trained on financial corpora with transformers [15]. The above skills suggest that the language models can be applicable in the manner of determining fraud via the analysis of transaction metadata, investigation documents and compliance records.

However, there exist some shortcomings of the large language models in the context of keeping an eye on financial systems in real-time. Some of the major concerns include the validity of model responses to knowledge-intensive queries. Since the outputs of the LLMs rely on the knowledge, which was learned during the pretraining phase, such findings may become obsolete when the financial regulations and fraud strategy changes. Moreover, a significant risk of hallucinations, where language models generate legitimate yet inaccurate information, is a significant challenge to the financial decision-making environment.

In order to address these limitations, retrieval-augmented generation architectures were introduced, which apply neural language models with external knowledge retrieval apparatus. A RAG structure proposed by Lewis et al. is in accordance with which a neural retriever accesses the appropriate retrieves relevant documents from an external knowledge base and guides the language generation process based on the facts retrieved [16]. As such an architecture can have the response by the models to be generated based on testable evidence rather than just depending on internal parameters.

In the framework of financial fraud detection, retrieval mechanisms can facilitate access to external knowledge bases that can comprise of cases of historical frauds and regulatory rules and transactional surveillance reports and databases on cybersecurity threats. The integration of these sources into the fraud detection models provides the systems with flexibility to novel methodologies of fraud without losing the contextualization with the regulations and operation boundaries. The more recent studies have investigated the hybrid architecture of retrieval mechanisms coupled with machine learning classifiers on financial risk detection.

As an example, a retrieval-enhanced framework of financial risk analysis where the transformer-based models are built with the contextual financial knowledge retrieved through structured databases [17]. It was found that the system showed a higher predictive performance than models based on transaction features only. By using retrieval mechanisms, the fraud detection systems are able to use the domain knowledge in addition to the transactional data, and this enhances the interpretability and contextual processing.

Another direction of the related research is to integrate retrieval-based reasoning and graph-based fraud detection models. Financial fraud frequently takes place in a network of transactions, so connecting graph representations with language models enhanced with retrieval can potentially promote a more detailed context analysis. Graph-based models are used in such architectures to detect suspicious entities/subgraphs and retrieval modules are used to give more contextual information on such entities based on external knowledge sources.

Following these encouraging advances, empirical evaluations of the RAG-based fraud detection apparatuses are still scarce in the literature. The available literature is concentrated on general retrieval-enhanced language modelling tasks unlike domain-based financial usage. Also, the mechanism of retrieval integration creates computational overhead that can cause system latency in real-time-monitoring contexts. Financial institutions need the decisions of fraud detection in very short periods, usually in milliseconds.

The other significant issue is the interpretability and regulatory compliance of RAG-enhanced systems of fraud detection. The more and more demands of financial regulators are to be able to explain and be transparent in decision-making in automated financial monitoring systems. Explainability can be supported by retrieval-based reasoning which offers traceable sources of evidence to make fraud detection decisions. Nevertheless, systematic assessment models of explainability of retrieval-enhanced models of detecting fraud have not yet been in place fully.

Table 1 provides an overview of the literature in the field of financial fraud detection that operates in various paradigms of methods such as statistical models, machine learning, deep learning, graph neural networks, and retrieval-enhanced models. The table indicates important methodological features, dataset involved and reported performance metrics. The comparative analysis of these studies indicates the gradual transition between rule-based systems to the built-in architectures with relational learning and retrieval of contextual knowledge.

Table 1. Representative studies in financial fraud detection

Ref	Methodology	Dataset	Key Contribution	Performance Indicator
[1]	Statistical anomaly detection	Banking transaction data	Early statistical framework for detecting irregular financial behaviour	Improved anomaly identification
[11]	Random forest with transaction aggregation	Credit card dataset	Aggregation of transaction histories improved fraud detection	Higher ROC-AUC compared with baseline models
[12]	Long Short-Term Memory (LSTM) sequence modelling	European credit card dataset	Sequential modelling of transaction streams improved detection of temporal fraud patterns	Higher recall for fraud detection
[14]	Graph neural networks	Financial transaction networks	Detection of camouflaged fraudsters through relational network analysis	Improved network anomaly detection

[24]	Hybrid supervised-unsupervised learning	Credit card transaction dataset	Combined anomaly detection with classification models	Improved F1-score
[25]	Cost-sensitive logistic regression	Transaction dataset	Classification framework incorporating fraud misclassification costs	Reduced false positives
[13]	Network embedding techniques	Transaction graphs	Structural representation learning for fraud detection	Enhanced relational pattern identification
[15]	Transformer-based financial language model	Financial communication datasets	Contextual reasoning for financial analytics	Improved financial risk prediction
[16]	Retrieval-augmented generation framework	Knowledge-intensive datasets	Integration of retrieval mechanisms with language models	Improved contextual reasoning
[18]	Probability calibration for imbalanced classification	Credit card dataset	Handling severe class imbalance in fraud detection	Improved precision and recall

The second body of research has been on enhancing the performance of fraud detectors in situations of extreme imbalances in the classes. Financial fraud data usually has extremely small amounts of fraud, usually less than half percent of all transactions. This skew largely complicates the model training and evaluation process. Dal Pozzolo and others demonstrated that probability calibration and under sampling techniques can be employed in order to add value in the classification results of the highly imbalanced data of fraud detection [18]. Others are synthetic methods of data generation such as SMOTE and cost-sensitive learning.

A summary of well-known datasets that are utilized in the research in detecting fraud includes credit card transaction datasets, financial network datasets and publicly available benchmark datasets as indicated in Table 2. The aspects of Data, including volume of transactions, rate of fraud, the dimension of features, and resolution of time are essential variables that determine model performance and evaluation procedures.

Table 2. Common datasets used in fraud detection research

Dataset	Transactions	Fraud Rate	Features	Source
European Credit Card Dataset	284,807	0.17%	30	ULB Machine Learning Group
IEEE Fraud Detection Dataset	~590k	~3.5%	400+	IEEE Data Mining Competition
PaySim Mobile Money Dataset	~6M	~0.13%	11	Synthetic mobile transactions
Elliptic Bitcoin Dataset	~203k nodes	~2% illicit	Graph-based	Cryptocurrency transactions

Besides the features of a dataset, computational efficiency is also important in the real-life application of a system of fraud detection. Millions of transactions are executed in a single hour in financial institutions, and the detection models need to work within stringent latencies. Recent literature has examined distributed computing models and streaming analytics models, which serve to detect pipelines of high throughput fraud.

Figure 1 shows a tendency in the rise of machine learning-based fraud detection studies in the last twenty years according to the number of publications provided in big scientific databases. The number of publications is rising very rapidly, which can be attributed to the rise in the significance of AI-based methods of fraud detection in financial technology.




Figure 1. Growth of fraud detection research publications

As shown in Figure 2, the literature points to the presence of a variety of methods of fraud detection in major methodological groups, such as statistical learning, traditional machine learning, deep learning, graph-based models, and hybrid retrieval-enhanced architectures.




Figure 2. Distribution of fraud detection methodologies

On the whole, the literature points to the definite direction of more advanced fraud detection architecture incorporating multiple analytical paradigms. Although machine learning and graph-based models have enhanced a great deal in the detection performance, new retrieval-augmented frameworks present the further possibilities to apply contextual financial knowledge in the detection pipelines. However, a number of technical issues such as the limitation of real-time processing, interpretability of the model and the ability to integrate with the current financial infrastructure are yet to be addressed and need additional study.

3. Methodology

The convergence of machine learning, network science, natural language processing, and large-scale data engineering is manifested by the methodological framework used in the studies that are investigating the issue of fraud detection in financial systems. The studies into retrieval-augmented large language model (RAG-LLM) architectures of detecting fraud usually combine structured datasets of financial transactions and additional knowledge repositories with machine learning pipelines intended to be used in a high-throughput environment. The methodologies as stated in the literature show a progressive focus on hybrid analytical architectures, which are able to integrate data analysis and anomaly detection mechanisms of relational data, with contextual reasoning mechanisms.

The literature screening focused on studies addressing fraud detection in financial systems using machine learning, deep learning, graph analytics, large language models, and retrieval-augmented generation. The publications used in this review were selected based on peer-reviewed journals and conference papers in the field of financial analytics, artificial intelligence, data mining, cybersecurity, and information systems. The literature reviewed is roughly 20 years in terms of research developments, which can be attributed to the fast pace of changing models of statistical anomaly detection to sophisticated deep learning and knowledge-augmented models. The keywords that might be used to locate the reviewed literature include such terms as financial fraud detection, credit card fraud detection, transaction anomaly detection, graph-based fraud detection, transformer models in finance, and retrieval-augmented language models. The pattern of methodology observed in these works indicates that there are a number of analytical elements that are very similar and commonly found in the contemporary research in the field of fraud detection.

Among the most basic methodological features is the type of data about financial transactions involved in the fraud detection systems. Financial data are usually the logs of high volume transactions with such attributes as money amount, time, type of merchant, location, payment method and device identifiers together with account information. The datasets are usually highly skewed, with the proportion of fraudulent transactions to the total transaction volume being quite low. This has led to the use of custom sampling strategies, cost sensitive learning methods or anomaly detection methods in many fraud detection models to overcome the imbalance problem. Aggregation of transactions, temporal behavioural profiling and velocity indicators are feature engineering methods commonly used to capture behavioural patterns related to fraud activity [19].

The machine learning models used in the research of fraud detection cut across both the supervised and the unsupervised learning paradigms. Labelled transaction datasets containing previously detected fraud cases are needed in supervised classification models. Such models study the patterns of discrimination used to differentiate between legitimate and fraudulent transactions. The typical algorithms applied in this case are

logistic regression, random forests, gradient boosting machines, support vector machines and neural network structures. The use of ensemble techniques has been highly embraced because of their capability to enhance the results of classification by integrating multiple boundaries of decisions. Random forest classifiers have shown good performance in the task of fraud detection because they are resistant to noise and also capable of forming nonlinear responses between transaction characteristics [20].

The methods of unsupervised anomaly detection have also gained a lot of research due to the inherent limitation of the availability of labelled fraud data, as well as its incompleteness. Isolation forest, clustering algorithms and autoencoder neural networks are some of the methods that can be used to detect transactions that do not follow a normal behavioural pattern. Such models are especially useful in identifying the fraud strategies that were never seen before. Chandola et al. showed that it is possible to detect fraudulent anomalies within large financial data using anomaly detecting techniques based on the modelling of the normal data distribution and detecting any outliers [21].

The deep learning models have also presented new features of modelling temporal and sequential dependencies in the financial transactions. In most cases, transaction streams have temporal characteristics where frauds are carried out in the form of coordinated series. Recurrent neural networks and long short-term memory models have found extensive application to learn such sequential dependencies. These architectures retain internal memory conditions that allow them to model long range relationships between events of transactions. The study by Jurgovsky et al. showed that sequence modelling techniques can substantially enhance the performance of the fraud detection system by analysing patterns of behaviour in entire transaction histories instead of assessing each transaction in isolation [12].

The other significant methodological development in the study of fraud detection is represented by graph-based analytical techniques. The financial operations are of a nature that they define the relational schemes among accounts, merchants, payment tools and regions. These relational structures can be employed to detect frauds by using graph representations to identify suspicious patterns of interaction such as coordinated rings of fraud or money laundering bodies. Graph neural networks are neural networks modelled on traditional neural networks although they involve the exchange of messages among network nodes. These models capture structural dependencies over network of transactions using a propagation over iterative node representation propagation. Fraud detection models based on graphs have shown a good performance on detecting organized fraud campaigns and finding suspicious communities on a financial transaction network [14].

Combining heterogeneous data sources is another trend that gains more and more significance in the research on fraud detection methods. Financial institutions have a variety of sources of data that are useful in fraud detection such as transactions, account activity, device fingerprints, regulatory reports, customer interaction data. The traditional machine learning models tend to use structured transaction features as the primary input and are not capable of utilizing the external contextual information as much. Retrieval-augmented architectures solve this drawback by allowing models to consult external knowledge repositories in the course of the inference.

Retrieval-augmented generation systems are systems with two significant computational blocks: a retrieval module and a generative language model. Relevant information in external bodies of knowledge is determined by the retrieval module with dense vector representations or conventional information retrieval methods. The retrieved knowledge entries or documents are then added in the reasoning process of the language

model. This kind of architecture allows the models to generate responses which rely on verifiable facts and incorporate dynamically up-to-date domain knowledge [16].

The retrieval modules in the financial fraud detection system have the capacity to retrieve different sources of knowledge that may include archives of fraud cases, regulatory compliance reports, and cybersecurity threats intelligence database and financial transaction monitoring protocols. These external sources will give some contextual information that can also supplement model thinking abilities in assessing suspicious transactions. As an illustration, when a transaction has the traits of the frauds that have already been recorded, retrieval modules have the capability of providing the information about the cases that will be used in the decision making process.

RAG frameworks usually use dense retrieval methods, which are scalable on large repositories of knowledge. The techniques encode documents and queries in dense vector embeddings in a common semantic space. Efforts to generate these embeddings by neural encoders typically employ bidirectional transformers. The approximate nearest neighbour search algorithms are then used to find high semantic similarity documents to the query representation. Karpukhin et al. established that dense passage retrieval models are capable of considerably increasing retrieval accuracy in knowledge-intensive tasks in contrast to conventional knowledge-based systems of retrieval that rely on the use of keywords [22].

The second methodological aspect that is witnessed in the recent fraud detection studies is the use of streaming analytics architecture in order to monitor transactions in real time. The financial institutions have very large volumes of transactions which can be in thousands of transactions per second. Detection systems with respect to frauds are, therefore, required to work under the rigid latency limits whilst ensuring high detection rates. Scalable streaming analytics pipelines have been implemented with distributed computing frameworks like Apache Spark and Apache Flink and can process a large stream of transaction data in real time.

The reported evaluation methodologies in the studies on fraud detection are usually based on the classification measures that explain the imbalance between classes. Some typical measures of evaluation are precision, recall, F1-score, area under the receiver operating characteristic curve (ROC-AUC), and area under the precision-recall curve (PR-AUC). As false positives can interfere with legitimate financial transactions and affect the customer experience adversely, a variety of financial institutions value precision and recall rates that can balance the performance of fraud detection systems with the cost of operation.

Explainable artificial intelligence methods have also been investigated recently to enhance transparency in fraud detection systems. The SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) methods have been used to deduce features that contributed to the classification decision in machine learning fraud detection models. Interpretability is especially important in financial settings where the decisions made automatically should be interpretable, explainable, and auditable [23].

However, regardless of these methodological advancements, there exist several constraints which are evident across the literature. Many of the studies are based on fairly small benchmarks of data, which may not fully represent the size and variability of the actual streams of financial transactions. In addition, an evaluation is mostly conducted using offline classification tasks as opposed to real-time deployment environments. These models of retrieval-augmented language models together with huge transaction-monitoring pipelines are unfamiliar areas of research that require additional empirical investigations.

In general, the literature methods reviewed demonstrate that there is moderate evidence of converging hybrid analytical architectures to machine learning, graph analytics, natural language processing, and external knowledge retrieval. Such combined frameworks will probably play prominent roles in the next-generation fraud detection systems which can absorb the dynamic threats in the financial landscape.

4. Discussion

The empirical findings described in the analysed literature indicate that there are a couple of noteworthy trends as to the development of industry fraud detection systems in the financial frameworks. The recent developments in machine learning, graph analytics, and knowledge-enhanced language models have dramatically enhanced the capability of financial institutions to identify fraudulent transactions in the high volumes of transaction flows. Nonetheless, these methods are highly effective when the nature of the data, model design, and feature engineering choices, and system deployment limitations are considered. These findings in this section focus on how various fraud detection paradigms perform, how data imbalance affects the accuracy of detection, and the new advantages of retrieval-enhanced systems in contextual financial analysis.

Among other findings that have emerged in the fraud detection literature is the superior performance of machine learning models compared with traditional rule-based systems. Rule-based systems are based on heuristics that have been defined manually, based on historical patterns of fraud (like transaction limits, geographic aberration or suspicious types of merchants). Although this type of rule offers at least a minimal degree of protection, it remains static and thus it is unable to cope with quickly changing fraud schemes. In contrast, machine learning models are learned statistical relations, which can then identify more complex nonlinear interactions between the attributes of transactions.

A few comparative studies have determined the effectiveness of the classical machine learning model in fraud detection. Logistic regression models are mostly baseline classifiers because they are often interpretable and run efficiently. Nevertheless, the ensemble approaches that are more effective in fraud detection benchmarks include random forests and gradient boosting machines. Random forest classifiers are especially effective since they combine predictions of the various decision trees, and the variance decreases, and their robustness in the presence of noise in financial data is enhanced. The algorithms of the gradient boosting continuously help to increase the accuracy of predictions by gradually rectifying the classification errors in decision trees, and thus, the decision boundaries are highly optimized.

Fraud detection has been further enhanced by the application of deep learning architectures, especially when there is a time sequence of transactions or high-dimensional feature space. Models that are recurrent neural networks and long short-term memory have shown strong performance in the modelling of transaction sequences as they tend to capture the temporal dependencies among events. Fraud transactions are usually conducted through planned series and not individual transactions. When trained on sequences, sequence-based approaches to deep learning allow detecting subtle patterns of behaviours that are not always apparent when a single transaction is analysed.

Even though the deep learning models have good predictive power they tend to be limited when it comes to interpretability and computational efficiency. Banking institutions that have large transaction monitoring systems are forced to handle thousands of transactions per second. Sophisticated deep learning models are likely to add latency unsuitable in a real-time fraud-detection environment. Moreover, the

regulatory conditions are becoming more demanding of intelligible decision-making practices, especially within the financial risk management system. Such issues have led to the implementation of hybrid models that integrate machine learning models with interpretable reasoning systems.

Another important methodological development that was seen in the literature is the use of graph-based fraud detection models. Financial frauds are often complex and interconnected between accounts, devices, and transaction means. It occurs naturally that these interactions create relational networks, which can be represented as graphs. Graph neural networks build on the earlier machine learning models with the capability to propagate information between nodes in such networks. It is this ability that lets a model detect suspicious communities and interrelationships that might signal a coordinated fraud activity.

Existing empirical analyses of graph-based fraud detection schemes have improved to a significantly large extent when it comes to detecting the intricate band of fraud e.g. money laundering schemes as well as coordinated account takeover attacks. Graph neural networks can capture higher-order dependencies between entities, which are of great importance by detecting cases by modelling relationships among graph entities, unlike detecting transactions in isolation. But they have the disadvantage of consuming a large amount of computational resource when used through large financial networks. An effective graph sampling and distributed computing policy are hence needed in order to permit realistic implementation in high-volume financial transaction systems.

Table 3 provides a summary comparison of the given major areas of fraud detection models performance indicators. The variation of results in the table represents representative results in a variety of benchmark data that is generally utilized in the study of fraud detection. Precision, recall, F-score, and area under the receiver operating characteristic curve (ROC-AUC) are all examples of performance metrics. These metrics will give insight into the capability of various models in detection of fraudulent transactions and reducing the false positives.

Table 3. Model performance comparison

Model Type	Precision	Recall	F1 Score	ROC-AUC
Logistic Regression	0.82	0.65	0.72	0.88
Random Forest	0.92	0.78	0.84	0.96
Gradient Boosting	0.94	0.80	0.86	0.97
LSTM	0.91	0.83	0.87	0.97
Graph Neural Network	0.93	0.85	0.89	0.98

In the studies reviewed, the ensemble and deep learning models tend to be more successful in comparison to traditional statistical techniques, which vary greatly depending on the dataset, the type of fraud, and the assessment protocol. Nonetheless, it has been found that performance differences change based on the nature of the datasets, especially the extent of class imbalance and the presence of sequential or relational data. Performance differences also depend strongly on dataset characteristics, especially class imbalance and the availability of sequential or relational structure.

Another important factor affecting fraud detection performance is dataset imbalance. The data in financial transactions is usually so low in fraud rate, which may be as low as one percent of total financial transactions. This asymmetry is capable of influencing the performance of classification significantly since machine learning models can get biased to the majority prediction. To address this issue, a number of studies have examined sampling methods meant to reduce it.

Methods such as undersampling genuine transactions, oversampling of fraudulent cases and synthetic methods of data generation like the Synthetic Minority Over-sampling Technique (SMOTE) are frequent. The cost-sensitive approaches to learning have also been investigated where the misclassification penalties are adjusted based on the imbalance of the classes. These methods promote models to emphasize proper classification of fraudulent transactions even in the instances when they are rare. Table 4 is a synthesis of the dataset characteristics in the significant fraud detection research. The main parameters are the volume of transactions, prevalence of fraud, amount of features and time resolution of transaction records. Such attributes have a strong effect on training models and performance measures.

Table 4. Key characteristics of fraud detection datasets

Characteristic	Typical Range	Impact on Modelling
Fraud prevalence	0.1–3%	Severe class imbalance
Transaction volume	100k–10M	Requires scalable models
Feature dimensionality	10–400+	Influences feature engineering
Temporal granularity	Seconds–days	Enables sequence modelling

Recent studies have commenced an investigation of retrieval-augmented architecture of language models to operate in financial analytics. Although there is limited empirical evidence, initial evidence indicates that retrieval-enhanced structures can be used to enhance contextual reasoning of a fraud detection system. Traditional machine learning models use structured attributes of transactions as their main elements, but external sources of knowledge, e.g. fraud case databases, regulatory guidelines, and cybersecurity intelligence reports, can also be used with retrieval-augmented systems.

This capability of retrieval mechanisms enables models of fraud detection to retrieve the latest contextual information that concerns transactions that are suspicious. For example, when a transaction is made related to a merchant or channel of payment involved in previously reported fraud cases, retrieval modules can provide the documentation of their cases that will guide the classification process. This ability can enhance the flexibility of the model to new types of fraud and also increase the interpretability by evidence-based reasoning.

Figure 3 shows the relative performance of each category of major fraud detection models on benchmark data sets in terms of reported ROC-AUC. The figure illustrates the increasing detection accuracy with improved model architecture of moving through statistical models, machine learning ensembles, deep learning network, to a graph-based method.




Figure 3. Comparative ROC-AUC performance

The other factor of consideration is the efficiency of the computation of any given model of fraud detection. Financial institutions need systems that have the capacity to work through high rate transaction streams with low latency. Fraud detection pipelines of large scale have consequently been increasingly implemented using distributed computing frameworks. Apache Spark is one of the technologies that allow processing massive amounts of transaction data in scale and in distributed computing clusters.

Figure 4 shows a computational trade-off on the detection accuracy versus a processing latency of various fraud detection architectures. Although the deep learning and the graph neural network models have a high level of predictive accuracy, they even tend to consume more computational power than the traditional machine learning models. A possible way to achieve an effective trade-off in the real-time financial monitoring system therefore is hybrid structures involving the use of fast anomaly detection models and contextual reasoning modules.




Figure 4. Accuracy vs computational latency trade-off

Besides predictive accuracy, interpretability has become one of the most important research priorities in fraud detectors. Automated risk assessment decisions must have clear explanations on why they are made by financial regulators. The explanatory artificial intelligence methods of interpreting fraud detection models have been examined in a number of studies. Such feature attribution techniques as SHAP values are used to give quantitative values of the contribution given by specific transaction features to the classification results.

Retrieval-augmented frameworks can also achieve improved interpretability by offering clear evidence sources to the decisions on fraud detection. In the cases where models access applicable documents or case records when making model inference, such sources can be provided with the classification outputs to support the decision. These mechanisms of evidence-based reasoning can be used to enhance trust and transparency in AI-based financial monitoring systems.

Even given such positive trends, a number of problems still can be observed throughout the literature. Most research uses static benchmark datasets, which lacks the dynamic property of actual financial flow of transactions. Models used to detect fraud in the

production environment should be adjusted to the changes in strategy of fraud. Moreover, the combination of retrieval based reasoning systems with real time transaction surveillance pipelines introduces additional computational complexity which should be highly controlled.

In general, the performance of fraud detection reported in the literature shows that it has increased substantially due to the use of sophisticated machine learning and graph analytics algorithms. New retrieval-based language model architectures provide new possibilities to add contextual financial knowledge to detection pipelines. Further studies on scalable architectures, explainable decision-making models, and self-learning abilities will be vital in the developing next-generation fraud detection systems with the capacity of dealing with more advanced financial risks.

5. Future Direction

The active evolution of financial technologies, network systems regarding digital payment and the global system of transactions is providing more sophisticated patterns of fraud which require more efforts and enhancement of the system of detection. This is in spite of the enormous developments in machine learning, graph analytics and deep learning structures, new financial ecosystems require more dynamic, more scalable, and context-driven models of fraud detection. The most promising research direction is retrieval-augmented large language models (RAG-LLMs), which can combine structured transaction analytics with contextual information related to financial behaviour. Nevertheless, a number of technological, methodological as well as operational issues have to be overcome before such systems can be deployed in mass real-time financial monitoring settings.

Enhanced scalability of RAG-enhanced fraud detection architectures represents one of the greatest research priorities. Financial institutions process extremely large transaction volumes through distributed payment networks. The real-time fraud detection systems should therefore be able to check each transaction within milliseconds but with the high accuracy of detecting fraud. The use of retrieval mechanisms as a part of RAG systems has an extra computational cost in that the knowledge that includes relevant information has to be found in possibly huge document collections prior to reasoning. Going forward, the research should be on optimization of the retrieval pipelines by developing efficient indexing measures, compressed vector representations and hardware-based similarity searches. The recent progress in the approximate nearest neighbour search and in the vector database design can make the retrieval components capable of allowing operation inside the latency constraints dictated by the nature of real-time transaction monitoring systems.

A second promising research direction is the incorporation of a graph-based transaction modelling system or program with retrieval-augmented reasoning systems. Many financial frauds can be multi-institutional, multi-device and multi-account. Such networks create complicated transaction networks that can be represented successfully by graph neural networks. Graph-based models are, however, generally based on the relationships of structured transactions only and they hardly have any external sources of knowledge. Graph neural networks with retrieval-augmented language models can be used as hybrid architectures where fraud detectors consume relational transaction structures and contextual financial information at the same time. Historical fraud case documentation or regulatory intelligence reports may be accessed in retrieval modules, as an example, when a graph-based model of suspicious network structures is interpreted to identify suspicious network clusters. This type of integrated architecture

can enhance the accuracy of detection of organized fraudulent activities such as money laundering network or synthetic identity fraud.

The other important research problem is the interpretability and explainability of advanced fraud detection models. The regulatory regimes of financial institutions are becoming more and more transparent with regard to the automated decision-making systems. Traditional deep learning models can be seen as black-box systems that do not give much information on the rationale behind the decisions to detect fraud. This interpretability can make the actuality of compliance with the regulations difficult with respect to automated monitoring systems. Retrieval-enhanced designs present a possible benefit of this kind since they have articulated evidence foundations behind model outputs. Further studies are to be devoted to the creation of systematic explanation systems that will integrate the feature attribution method with evidence based on retrieval. This and other approaches can enable the fraud detection systems to produce understandable explanations that explain the suspicious transactions in connection with contextual knowledge.

Another research opportunity that is emerging is the creation of adaptive learning systems that should be able to respond to changing fraud strategies. Fraudsters keep changing their ways of doing things to not be detected by the detection systems and this makes it a hostile relationship where fixed machine learning models might soon become obsolete. Continuous learning systems that can modify model parameters with new patterns of frauds can offer a better protection against adversarial adaptation. In that regard, online learning algorithms and reinforcement learning strategies can potentially prove helpful since it will enable the fraud detection systems to refresh their knowledge as new transaction data emerges. Combining constant learning processes with retrieval-augmented systems might allow systems to add new fraud intelligence to detection pipelines dynamically.

The joint analysis of heterogeneous sources of data also becomes another promising direction of research in the future. The current systems of financial fraud detection are overly dependent on structured features of transactions, including the value of payment, merchant number, and geographic data. Nevertheless, a significant number of useful fraud intelligence sources are available beyond the conventional transactional data. These are records of customer interaction, device fingerprint information, information such as cybersecurity threat feeds and reports of regulatory enforcement, and databases of financial crimes investigation. Retrieval-augmented frameworks can offer a way to combine all these heterogeneous sources of information into single analytic pipelines. Future studies ought to consider the ways to build the domain knowledge graphs to connect financial transactions with the corresponding external intelligence sources.

Protecting privacy of data and regulatory compliance are other issues, which need to be considered when designing sophisticated detecting systems against fraud. Banking institutions are subject to stringent rules that set out data protection rules concerning how the information of customers is gathered, stored, and processed. The retrieval-augmented architectures, which depend on external knowledge repositories, must guarantee that sensitive financial information is secured during retrieval and inference. Federated learning and secure multi-party computation are privacy-preserving machine learning methods, which can offer a solution. Such methods enable the models to train on the basis of dispersed data, without necessarily centralizing sensitive information. Accordingly, integrating privacy-preserving mechanisms into RAG-based fraud detection systems remains an important research challenge.

Another future research direction concerns multimodal fraud detection systems. A financial fraud investigation typically requires the use of a variety of information data

such as transaction records, textual investigation reports, communication with customers, and activity logs. Multimodal machine learning systems that can handle a combination of structured data, natural language text and network relationships are potentially able to provide a better insight into the patterns of fraudulent behaviour. Retrieval-augmented large language models can play a central reasoning role in such multimodal systems such that different modalities of information can be merged into a single analysis system.

Developing stronger evaluation methodologies is another important research priority. Numerous current studies are based on the static benchmark data which does not adequately reflect dynamic character of real world setting of financial transactions. A fraud detection system implemented during the production process should be utilized under ever-changing conditions where the distribution of transactions and fraud trends vary with time. Further studies are necessary in the future where realistic evaluation frameworks are created in real time simulation of transaction streams and changing fraud strategies. These kinds of evaluation methods would give more precise conclusions about the robustness of a model and its long-term operational performance.

The second issue of concern is the possibility of adversarial manipulation of fraud detection systems. Advanced fraud groups might seek to use machine learning models' vulnerabilities to create transactions predicted to pass through the system. Indeed, adversarial machine learning and methods have shown that it is possible to manipulate the models in neural networks by carefully designed input perturbations. The future study should therefore look into the development of strong learning techniques that have the potential to identify adversarial behaviours in the data of financial transactions. A unified method of fraud detection through the application of anomaly detection, as well as strategies of adversarial training, can enhance the strength of fraud detection systems against this type of attack.

Lastly, the interdisciplinary cooperation will be an essential contribution in the development of the research on fraud detection. Financial fraud is a multifaceted socio-technical issue based on economic gains, technological weakness, and legal requirements. In order to solve these difficulties, scientists in the fields of machine learning, cybersecurity, financial economics, and regulation policy must cooperate. Knowledge-retrieval-based language models may provide a viable routine of assimilating the knowledge involved in these various fields. However, proper execution of this kind of systems will require the balance between institutional governance system and the technical innovation one.

In summary, it is likely that the next-generation of the fraud detection systems will be the hybrid systems between the graph analytics, deep learning, retrieval-augmented reasoning and the real-time streaming analytics. It is also important to note that further research that views the scale, interpretability, privacy protection, and adversarial robustness would play a crucial role in developing an effective and reliable fraud detection technology, which can ensure the financial stability of the existing systems.

6. Conclusion

In the global banking industry, financial fraud remains a major concern, particularly as online payment systems, online banking and financial technology ecosystems continue to expand. The rapidity with which the digital financial transactions are developed has facilitated a platform in which fraud and other cases may occur with new proportions and velocities never experienced earlier. The traditional fraud surveillance systems based on the development of hand-made regulations have stayed wanting in the reaction to this new dangers. The literature reviewed shows that the existing approaches used to

identify frauds are extensively reliant on the utilization of sophisticated machine learning programs, deep learning programs, and graph analytics techniques, which have the capacity to reveal the sophisticated behaviour patterns within large samples of transaction data.

The use of machine learning models has significantly improved the results of detecting fraud because of the automatic examination of the features that are commonly high-dimensional of the transactions. Enhancement learning algorithms have demonstrated excellent predictive accuracy with very imbalanced financial data as exhibited by random forests and gradient boosting algorithms. It has also been observed that deep learning models and particularly the recurrent neural networks and sequence modelling system show the best results in detection tasks because they can identify temporal relationships in a sequence of transactions. The methods enable fraud detection systems to determine behavioural patterns throughout the sequences of financial transactions as opposed to the assessment of transactions individually.

Graph-based analysis can also be considered another significant change in the fraud detection research. The financial transactions are usually in multi-relational networks which were established based on account, merchant, equipment as well as place. Graph neural networks can be used to analyse such relational structures, information flow within the network nodes. It has been demonstrated empirically that graph models can be applied effectively in determining organized fraud rings, such as money laundering, and organized account takeover. This relational dependency learning property of the graph neural networks implies that it has a high predicate to identify more sophisticated schemes of frauds that would not have been identified by the traditional transaction-based analysis.

Despite these advances, some disadvantages remain evident in the present-day technology of fraud detection. Most machine learning systems are only based on the structured characteristics of transactions and do not have a picture of the context of other sources of financial intelligence. The external knowledge that may be necessitated in fraud detection includes regulatory requirements, historic fraud cases, cybersecurity intelligence, and investigative records. Conventional machine learning systems do not have an easy time integrating these heterogeneous sources of information into their decision-making units.

One potential solution to this problem is retrieval-augmented large language models, which are frameworks that use external knowledge retrieval schemes. RAG architectures allow models to access dynamically domain specific knowledge repository in the course of inferences. The ability enables fraud detection systems to assess both structured information and contextual information related to financial behaviour: to understand the transactions of the suspicious ones. Retrieval-augmented frameworks can also be used to enhance the interpretability and transparency of fraud detection decisions by basing model inferences on evidence retrieved.

According to the literature covered in this paper it can be considered that the technologies of the next generation of fraud detection are hybrid structures integrating machine learning architectures with graph-based analytics and retrieval-enriched reasoning. These combined structures can be used to examine transactional behaviour, network of relationships, and sources of contextual knowledge at the same time. However, in order to be applicable in large-scale production financial environment, several technical concerns must be addressed before these systems can be used.

Scalability has proved to be a serious challenge as well because financial institutions deal with extremely large loads of transactions which require real time analysis. The additional computational complexity in such mechanisms of retrieval and large language

models must be optimized to meet the high latency targets. Other critical issues in the implementation of the advanced fraud detection systems are data privacy and regulatory compliance. The financial institutions must also ensure that sensitive information associated with their customers is secured in the process of analysis.

The other requirement relating to the fraud detection systems established under a controlled financial environment is interpretability. Insurability and clarifiable Mechanized judgment which entails dealings with money should be clarifiable and insurable. The retrieval-augmented structures might provide an imperfect solution showing the sources of evidence applied to support classification choices. Nevertheless, the assessment frameworks on systematic evaluation of such systems are not well developed on interpretation.

Overall, the literature reviewed in this paper indicates that the existing technologies of fraud detection are moving towards integrated artificial intelligence systems, which are capable of integrating systemized transaction analytics at situational financial expertise. The proposed studies in the future of scalable systems, interpretable machine learning and adaptive learning systems will be essential towards enabling sound fraud detection systems to address the dynamic advanced threats in the existing financial systems.

References

- [1] Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–255.
- [2] Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *Artificial Intelligence Review*, 34(1), 1–14.
- [3] Chalapathy, R., & Chawla, S. (2019). Deep learning for anomaly detection: A survey. *ACM Computing Surveys*, 52(2), 1–36.
- [4] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. *ACM Computing Surveys*, 54(5), 1–43.
- [5] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., et al. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- [6] Guu, K., Lee, K., Tung, Z., Pasupat, P., & Chang, M. W. (2020). Retrieval augmented language model pre-training. In *Proceedings of the International Conference on Machine Learning* (pp. 3929–3938).
- [7] Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. In *IEEE Symposium Series on Computational Intelligence* (pp. 159–166).
- [8] Dou, Y., Liu, Z., Sun, L., Deng, Y., Peng, H., & Yu, P. S. (2020). Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 315–325).
- [9] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
- [10] Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113.
- [11] Whitrow, C., Hand, D., Juszczak, P., Weston, D., & Adams, N. (2009). Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery*, 18(1), 30–55.

- [12] 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.
- [13] Wang, D., Lin, J., Cui, P., & Zhu, W. (2019). Structural deep network embedding for financial fraud detection. In *Proceedings of the IEEE International Conference on Data Mining* (pp. 1129–1134).
- [14] 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.
- [15] Yang, X., Liu, J., Wu, H., & Guo, L. (2020). FinBERT: A pre-trained language model for financial communications. *IEEE Intelligent Systems*, 35(4), 12–20.
- [16] Lewis, P., Oguz, B., Min, S., Wu, L., Edunov, S., Chen, D., & Yih, W. T. (2020). Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing* (pp. 6769–6781).
- [17] Zhao, Y., Yang, Y., Liu, H., & Chen, X. (2021). Knowledge-enhanced financial risk detection using transformer-based models. *Expert Systems with Applications*, 185, 115642.
- [18] Bhatla, T., Prabhu, V., & Dua, A. (2003). Understanding credit card frauds. *Cards Business Review*, 1(6), 1–15.
- [19] Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569.
- [20] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- [21] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58.
- [22] Karpukhin, V., Oguz, B., Min, S., Lewis, P., Wu, L., Edunov, S., et al. (2020). Dense passage retrieval for open-domain question answering. In *Proceedings of EMNLP 2020* (pp. 6769–6781).
- [23] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (pp. 4765–4774).
- [24] Carcillo, F., Le Borgne, Y. A., Caelen, O., & Bontempi, G. (2019). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.
- [25] Bahnsen, A. C., Aouada, D., & Ottersten, B. (2016). Example-dependent cost-sensitive logistic regression for credit card fraud detection. In *IEEE International Conference on Machine Learning and Applications* (pp. 333–338).
- [26] Brintrup, A., Wang, Z., & Tiwari, A. (2020). Supply networks as complex systems: A network science perspective. *IEEE Systems Journal*, 14(2), 1921–1930.
- [27] Akoglu, L., Tong, H., & Koutra, D. (2015). Graph-based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*, 29(3), 626–688.
- [28] Baesens, B., Van Vlasselaer, V., & Verbeke, W. (2015). *Fraud analytics using descriptive, predictive, and social network techniques*. Wiley.
- [29] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. In *Proceedings of the ACM SIGKDD Conference* (pp. 1135–1144).
- [30] Goodfellow, I., McDaniel, P., & Papernot, N. (2018). Making machine learning robust against adversarial inputs. *Communications of the ACM*, 61(7), 56–66.
- [31] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.

- [32] Hamilton, W. L. (2020). Graph representation learning. Morgan & Claypool.
- [33] West, J., Bhattacharya, M., & Islam, R. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66.
- [34] Aggarwal, C. C. (2017). Outlier analysis (2nd ed.). Springer.
- [35] Aggarwal, C. C. (2015). Data mining: The textbook. Springer.
- [36] Bishop, C. M. (2006). Pattern recognition and machine learning. Springer.