Deep Science Publishing, 2025
https://doi.org/10.70593/978-93-7185-061-2_8

DeepScience
Open Access Books

Chapter 8: Real-Time Fraud Detection through Agentic Systems

8.1. Introduction

Agentic Systems for Real-Time Fraud Detection investigates a real-time fraud detection approach that combines agentic systems with real-time data processing. Agentic systems are capable of autonomous, goal-directed activity and—when enabled to make decisions independently—they can react to suspicious activities quickly without human intervention. By detecting fraudulent transactions at first sight, the risk of financial loss induced by fraudulent transactions is minimized. The sections “Real-Time Data Processing” and “Fraud Detection Algorithms” identify challenges related to the real-time collection and streaming of data, outline techniques that support real-time data processing, and discuss methods for recognizing fraudulent transactions. The principal motivation for a real-time fraud detection approach is explained in “Background and Motivation.”

The demand for real-time fraud detection derives from the fact that the faster a fraudster is detected, the smaller the amount of money lost—hence, the greater the incentive for financial institutions to implement real-time fraud detection. Given that large financial institutions generate thousands (or even millions) of transactions daily, not all suspicious activities can be investigated by humans; therefore, automating the credit card risk management process becomes imperative. Service providers typically have access to cardholders' historical transaction details, ratings, and reports from various resources. When a new transaction is initiated, the service provider assigns a risk score, which helps the cardholder monitor the state of transactions in real time. Nevertheless, service providers’ decisions depend on how the fraud detection models are built—specifically, the methods of fraud and risk score calculations and the nature of the data on which the models are built, such as whether data are collected in a batch or streaming mode. An outline of these challenges is provided in “Future Directions.”

8.1.1. Overview of the Research Framework

Fraud detection remains an ongoing challenge in many industries because it prohibits economic development, damages brand reputation, and can lead to legal penalties. Detecting fraud immediately can protect institutions and act as a deterrent while minimizing the impact of specific cases. A research framework is proposed to investigate real-time fraud detection through agentic systems, which are capable of autonomous task completion.

The framework finds its need in the high costs of fraud in the digital age and the emergence of agentic systems. Evidence of the potential efficacy of agentic systems is found with companies combining real-time data streaming and advanced algorithms for fraud or anomaly detection, whereas published research has largely focused on post-fraud analysis. Real-time data processing is integral to the fraud detection process, and the framework accommodates that through a modular approach.

8.2. Background and Motivation

Real-time fraud detection is an increasingly important field as financial activities and commerce move to electronic and online platforms. Advances in communications networks and computer processing have created large corporate infrastructures for collecting vast quantities of customer transaction and other related data, resulting in “Big Data.” Although this data is valuable, manually reviewing it for signs of fraud and other illegal activities is impractical given both the volume of data and that the data is captured over extended time periods. Successful fraud detection depends on highly sophisticated computer processing so that a large amount of incoming data can be scanned continuously to trigger an alert, while rejecting legitimate activity that is unusual but not fraudulent [1-3].

Trying to detect fraud after the data is stored is not an optimal approach. The ability to analyze transactions as they occur is important because it gives the business the earliest possible notice to prevent continued fraudulent activity or even stop the current fraud transaction. For example, imagine that a user’s credit card has been stolen and used to buy 10 video games online. The credit card company can be notified of fraudulent activity only after the first transaction leaves their system. During the time that CoB transactions or settlement are processed, criminals can continue an attack on the same account and purchase thousands of dollars in illegal merchandise or even withdraw money from the user’s bank account.

8.2.1. Rationale for the Study and Its Importance

The rapid development of information and communication technologies has radically influenced people's daily lives, including how financial transactions are carried out. Mobile devices have enabled many people to operate banking applications anywhere and at any time. Given the millions of transactions that occur every second, financial organizations face the challenge of identifying fraudulent activities as quickly and efficiently as possible. While many organizations have internal procedures to prevent fraud, most still need to generate alerts or notifications in real time. The simultaneous growth in the use of social network applications, the increasing demand for personal identification data, and the development of banking devices were exploited by hackers to commit fraud and scams. Consequently, it is imperative that data be collected in real time to detect these kinds of activities and undertake the necessary control actions.

Fig 8 . 1 : Real-Time Fraud Detection: An Agentic Systems Approach

Thus, the rationale for this study is the need for systems that automate fraud detection and maintain constant interaction with real-time data, thereby making operations more agile and efficient. The objective is to explore the integration of agentic systems with real-time fraud detection, examining both current applications and emerging trends in the field. Given agentic systems' ability to operate independently and make decisions such as accomplishing assigned tasks, analyzing information, and responding to service

requests, their application in the development of fraud detection mechanisms is particularly relevant. This exploration establishes a framework within the broader domain of real-time fraud detection.

8.3. Literature Review

The literature on fraud detection is extensive, with a significant focus on financial applications. Although real-time fraud detection has received considerable attention over the years, agentic systems in this context are relatively unexplored. Nevertheless, the literature associated with agentic systems provides a foundation for understanding their potential role in real-time fraud detection.

Historically, the industry has prioritized offline fraud detection, gathering data and analyzing it post-transaction. As real-time data streaming infrastructure evolved, some researchers developed machine learning-based real-time fraud scoring systems. However, these systems generally operated on the backend, with human judgment ultimately guiding actionable decisions. Detecting and acting on fraudulent behavior necessitates screening data as it is streamed, analyzing the risk, and enabling the system to autonomously take action. More recent efforts consider the application of agentic systems to real-time fraud detection. An agentic system is defined as one that, in response to changing environments, assesses situations, formulates plans, and takes actions. Such systems are categorized in various ways, with one taxonomy identifying three types: assistant agents, executive agents, and mediating agents.

8.3.1. Historical Context of Fraud Detection

The history of fraud detection is as old as fraud itself. At the start, forensic examination of physical evidence was the main tool. Increasingly advanced methods were developed, including fingerprint analysis, DNA testing, and behavior and Cognitive Identity Profiling (CIP). Rapid industrialization in the nineteenth century brought new opportunities for fraud. As financial fraud increased, forensic accounting was developed. Now, in the twenty-first century, information technology made the global community more vulnerable to fraudulent transactions. Increasing numbers of transactions can be carried out online in real time; unauthorized users can also access the system.

8.3.2. Current Trends in Fraud Detection Technologies

The needs of the financial sector, a key driver for the introduction of different technological advancements for real-time fraud detection, also guided early efforts of

research [1,3-4]. Works such as that of Dodier and Mays that combines Hidden Markov Models (HMMs), supervised learning and unsupervised learning for pattern recognition, clustering and detection of anomalous user's behaviour represent a step further towards an agentic systems approach to fraud detection.

In parallel, the HMM approach was also combined with statistical rules for discovering both behaviour that deviated from the norm and behaviour that could be classified as fraudulent. The availability of consumption data from mobile telephony providers created demand for new real-time data-handling systems and algorithms capable of detecting fraud on one hand, and handling large streaming data sets on the other. This challenge contributed to the development of the Fraud Detection Repository (FDR), a system that integrates the fraud detection algorithms described above with a set of streaming algorithms.

8.3.3. Agentic Systems Overview

Agentic systems are designed to exhibit a form of agency, characterized by autonomy, proactive behavior, and goal-directedness. They can analyze their environments and

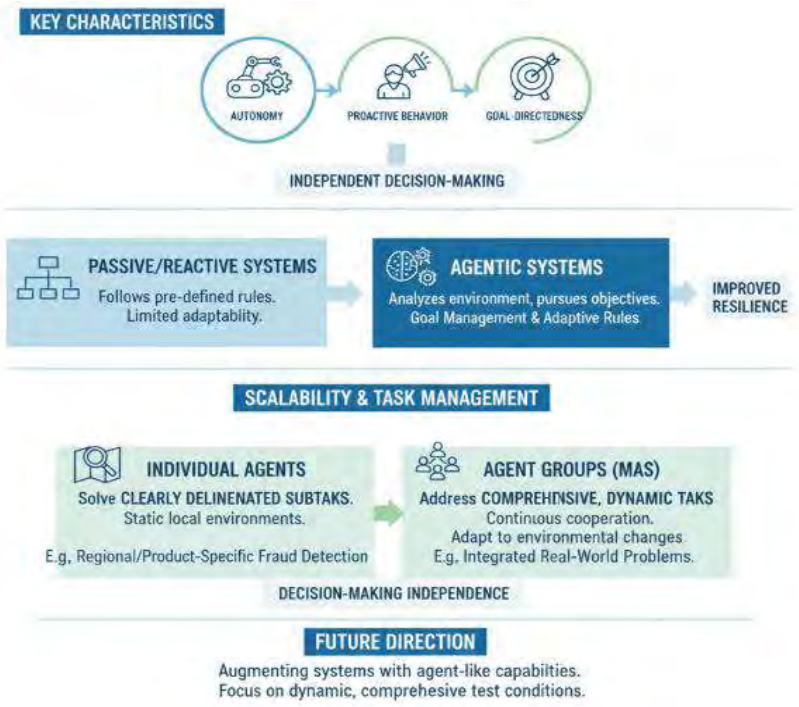

Fig 8 . 2 : Agentic Systems: Architecture & Applications

deploy actions in pursuit of objectives without needing detailed instructions for every step. Agentic systems are distinguished from passive or reactive systems by their capacity for independent decision-making. Many systems that currently operate following well-defined rules could be augmented with agent-like capabilities such as goal management and adaptive rule generation. These additional abilities may also improve the resilience of non-agentic systems, especially in complex Multi-Agent Systems (MAS).

One primary distinguishing feature among different agentic systems is the extent of their decision-making independence, which influences their scalability within MAS. Individual agents typically solve clearly delineated subtasks—static subsets of the overall environment—rather than the complete problem. Static subtasks provide relatively stable local environments over time, shaped but not solely defined by the agent's own actions. For instance, in fraud detection, an agent might oversee a regional or product-specific area. Conversely, agent groups address more comprehensive, dynamic tasks, requiring continuous cooperation to adapt to environmental changes. Real-world challenges often emphasize dynamic, comprehensive test conditions, reflecting the integrated nature of many practical problems.

8.4. Agentic Systems in Depth

Agentic Systems lie at the heart of every digital system. Scarcely one "autonomous" system could assert full autonomy, remaining dependent on external data collection and result verification. This dependence delineates the essence of the discussed Agentic Systems. Agentic Systems act decisively, fulfilling assigned mandates—provided they possess the autonomy to access the requisite real-world data or, conversely, data is actively fed into them. In the absence of such real-time data, Agentic Systems lack significant advantage.

A system must first ascertain whether a detected transaction constitutes financial fraud before activating an agent in response. This initial detection is a classic pattern recognition problem [3-5]. Only systems capable of issuing orders, rather than awaiting external commands, truly embody Agency. Moreover, the measurement of a transaction's potential for fraud must occur in real-time. Prolonged processing periods undercut the Agentic nature of the system, as might be presumed.

8.4.1. Definition and Characteristics

Agentic systems are software systems characterized by delegated, goal-directed decision-making under conditions that are—and may remain—uncertain. History has brought technology to a point where certain straightforward, well-circumscribed, repetitive tasks related to decision-making, deliberation, and execution can be delegated to or assigned AI technologies. One of the early tasks well suited to automation was the flagging of fraud and anomalies—from credit card use to e-commerce, financial transactions to election results. The effectiveness exhibited by earliest-generation credit-card fraud mitigation systems vindicates the concept. However, existing deployed systems suffer from a common flaw: transactional data is analyzed only after the transaction is completed, once the potential to prevent the fraudulent transaction has passed. This shortcoming is rooted in the way in which data is collected, created, and stored, and it can only be overcome through dynamic agentic collaboration.

Mainstream systems no longer have access only to erasable combinations of transactional data. Extant, as well as emerging, technology offers an unprecedented opportunity to observe—and to judge the ordinariness (commonality) of—transactional data at the moment the transaction occurs. Potentially suspected for the purposive fraud risk they represent, transactions can be put on hold for additional assessment. Recent advances in streaming ingestion technology make it possible to create agentic detection systems that determine membership in high-likelihood set(s) before the actual authorization for the transaction takes place. Client systems permit verification of the flagged transaction through category-specific queries, including biometric information, for example.

8.4.2. Types of Agentic Systems

The term agentic system can be interpreted in different ways. In a broad sense, an agentic system is a system that exhibits agency. Agency may be defined as the capacity for autonomous action [1,2-4]. In the context of human–computer interaction, the theory of agentic communication establishes that people interact with computing artefacts as not only recipients of information but also sources of social responses. Specifically, the theory of agentic communication asserts that when humans perceive a medium as agentic, they ascribe at least a modicum of autonomous actions and intentions to the medium and interact with it accordingly. An agentic system is therefore here defined as a system that people perceive as exhibiting agency.

The narrower view of artificial agent agency, rooted in verbal and virtual agent research, can be considered a special case of the broader definition. The theory of reasoned action posits that individuals form behavioural intentions based on attitude, control and

normative factors. Expanding on this theory, the theory of agentic action suggests that agentic actions in an interaction with a software agent are akin to social influence attempts by the agent. Social influence interpersonal concepts can be divided into categories of attitude change, compliance and acceptance. Behavioural intention classifications correspondingly address behavioural intention, behavioural compliance and behavioural acceptance. These concepts ground a model of agentic actions, delineate two-way agentic communication and specify recommendations for investigating agentic action etiquette.

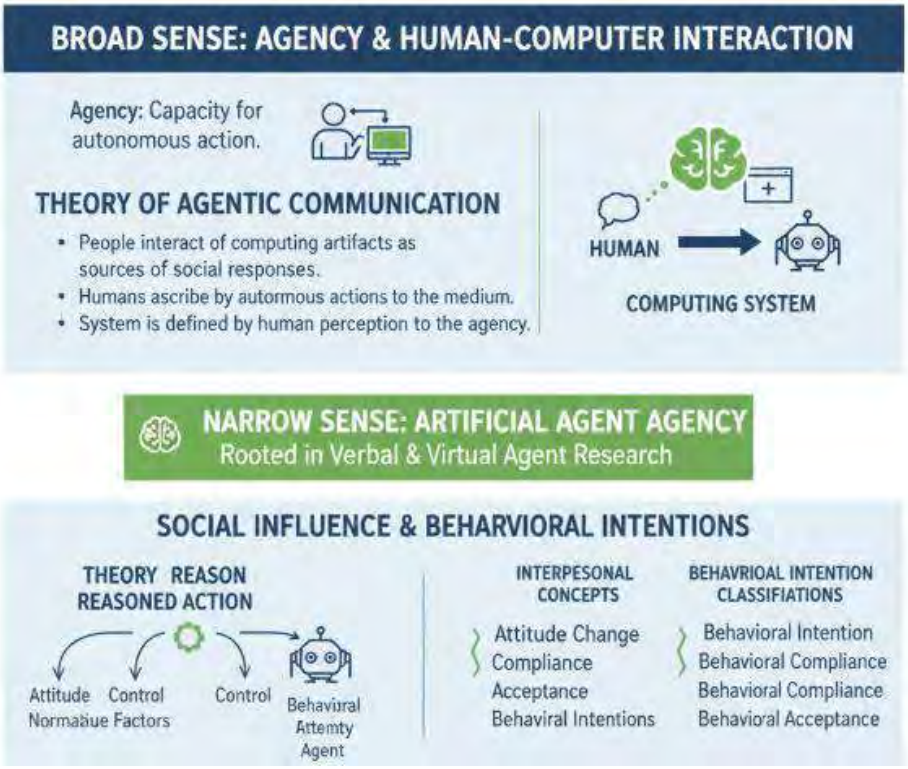

Fig 8 . 3 : Defining the Agentic System: Broad vs. Narrow Views

8.4.3. Applications in Various Domains

In everyday life, agentic systems play an increasingly important role in diverse domains, such as assistance and security. As examples, product recommender systems strive to propose products that satisfy existing customer wishes and needs, while process optimization systems continuously analyze product and machine data to reduce production costs and enhance product quality.

Fraud detection systems constitute a special category of agentic applications. The principal objective of fraud detection is to shield customers and companies from fraudulent acts. Fraud is defined as an act designed explicitly to cheat or deceive others, aiming to secure unlawful gains or advantages.

In the field of information processing, fraud encompasses any action contravening information usage guidelines that results in unauthorized, yet potentially permitted, use of computer systems or networks. Other interpretations stress the necessity of unlawful profit for the criminal, achieved by deceit, concealment, or a false statement of fact.

According to the U.S. Government Accountability Office, fraud detection denotes a process capable of preventing or detecting fraud, misappropriation, or malfeasance.

8.5. Real-Time Data Processing

Real-time fraud detection demands tools and methods that can keep pace with financial transactions on a split-second basis. Consider the example of credit card fraud: a decade ago, security teams combed through transactions months or even days later to identify the tell-tale signs. By the time they did, the bad actors in question often had emptied accounts, accumulated large charges, and moved onto their next target. Today, however, an agentic system can take in the data stream, detect the chicanery, halt the transactions, freeze the accounts, and notify all parties within moments. Running in real time thus requires moving beyond merely identifying fraud: the system must take the required preventative action too.

Many industries have a need for real-time data analysis. Self-driving cars respond to immediate information about road conditions. Health applications can alert users to fluctuations in vital signs. Moving beyond analysis, however, requires an agentic system. An agentic system simply has the authority to make decisions—initiating autonomous responses to a given event rather than halting to request approval beforehand. Self-driving cars can change lanes to avoid an accident. Health applications can dispatch emergency personnel. Fraud detection can immediately shut down a card or a transfer. Although reactive frameworks that signal the user to carry out the suggested actions have existed for some time, only proactive mechanism have the impact and scale required to provide true confidence in the system.

8.5.1. Data Acquisition Techniques

Real-time fraud detection hinges largely on the availability of data and the ability to process it continuously and in transmissions of low latency. The data is collected, aggregated, and processed by an adequate sensor network that can be a combination of

modern technologies like the Internet of Things and crowd sensing. Several types of transaction-related data need to be acquired and processed such as terminal, transactional, financial, and identity data. The integrity, privacy, and legality of the processing phase must be assured by the considerations made during the acquisition. While merely detecting fraud is already a complex topic due to the nature of the domain, attempting immediate mitigation introduces a further set of complications. It is essential that the intervention does not provoke inconvenience for the customer.

8.5.2. Data Streaming and Processing Frameworks

Real-time fraud detection requires certain data processing characteristics that are not prevalent in traditional analytics workflows. These include robust architectures for streaming, real-time scoring, and real-time processing. The Iris platform fulfills these requirements by incorporating several commonly used streaming frameworks, with Apache Kafka serving as the message queuing system that ensures scalable and extendible data transportation.

Apache Spark operates in a Kafka consumer role, performing analytical computations and fraud classification on individual messages. Spark can be configured to handle streaming data from Kafka in two modes: micro-batch processing for quasi real-time interaction or message-by-message interaction representing true real-time processing. Spark's analysis outputs, as well as categorized warnings and alerts raised by the Intelligent Engine, are streamed back into Kafka. From there, alerts and warnings are forwarded through the Iris platform to the Intelligent Engine and onward to Storbec for displaying in the user interface [6-8]. Special attention is given to Athena and Cassandra, which are used as sink points for Spark's results, making the analytical outcomes accessible for interactive user query and exploration.

8.5.3. Challenges in Real-Time Data Handling

One of the main difficulties in real-time data processing lies in integrating and preparing raw data for downstream analysis. Data are often collected from various sources and formats, even within the same organization and application. Such heterogeneous data require a flexible underlying infrastructure. Many end-point devices can provide data, including Web, mobile, Internet-of-Things, social, and enterprise applications. To capture all relevant information, the data platform should be able to process disparate datasets in various formats, such as structured (e.g., relational data), semi-structured (e.g., JSON), or unstructured (e.g., logs). Beyond the variety problem, the V's of Big Data—Volume, Velocity, Veracity, and Value—must also be addressed.

8.6. Fraud Detection Algorithms

The detection of fraud carries both an operational and a financial risk and as such it is important that fraud is detected as quickly as possible. Detection algorithms can be categorised into machine learning approaches, statistical approaches and hybrid approaches that combine the two. Within each category is a variety of algorithms that are suited to a specific type of detection problem such as supervised, semi-supervised, unsupervised, batch or real-time.

Agentic systems deployed in the context of real-time fraud detection must be capable of decision-making which impacts the live transactional flow and, as such, must operate with a level of certainty and confidence which would often require some level of supervision before deployment. The level of supervision that a model requires will govern the human interaction that is deemed necessary and informs the trust decisions in the interaction model. On this basis, currently available fraud detection algorithms tend to have a semisupervised or supervised approach and are typified by the categorisation in Table 1.:

8.6.1. Machine Learning Approaches

Automated decision-making approaches for credit-card fraud detection rely on machine learning techniques. A data-analysis model is trained on a dataset of processed features and their corresponding labels (fraud or legitimate), with the main goal of correctly assigning the labels to new, unseen data instances. Standard machine-learning algorithms often include neural networks, support-vector machines, decision trees, random forests, and k-nearest neighbors.

Every proposed method has advantages and disadvantages. For example, decision trees provide results that are easy to interpret and explain autonomously detected fraud classes; however, their automatic nature can make them too reactive [1,5,6]. On the other hand, artificial neural networks can detect complex correlations in the data, but their decision-making process requires expert human judgment. Supervised classification algorithms require a labeled dataset for training, which is not always readily available. Unsupervised learning methods such as clustering and entropy-based approaches have also been explored. Since data-driven approaches depend heavily on the dataset, techniques like k-fold cross-validation are used to ensure robust performance estimates. For performance comparison, measures including accuracy, precision, recall, false positive rate, and F-measure are employed.

8.6.2. Statistical Methods

Statistical methods represent one of the traditional approaches to data analysis. Some commonly used techniques are regression analysis, Bayesian estimation, Markov models, and Probit and Logit models. A characteristic of these methods is that they usually work well when there is abundant data. For example, they make the best possible prediction in a single point based on the previously known data. However, they do not perform well in dynamic situations when new data become constantly available, as it occurs in real-time systems.

In the context of classification, statistical methods require aggregating data into equal-width bins with predefined borders based on expert knowledge, implying that all values from one bin are equal. Then, the analysis is performed on these bins instead of analyzing every data item. Furthermore, the selection of these borders requires domain expertise that is often unavailable. Finally, real-time detection intensity, which keeps adapting to dynamic changes, is not yet addressed. Recent developments in data mining and machine learning partially overcome these limitations by generating patterns from the experience.

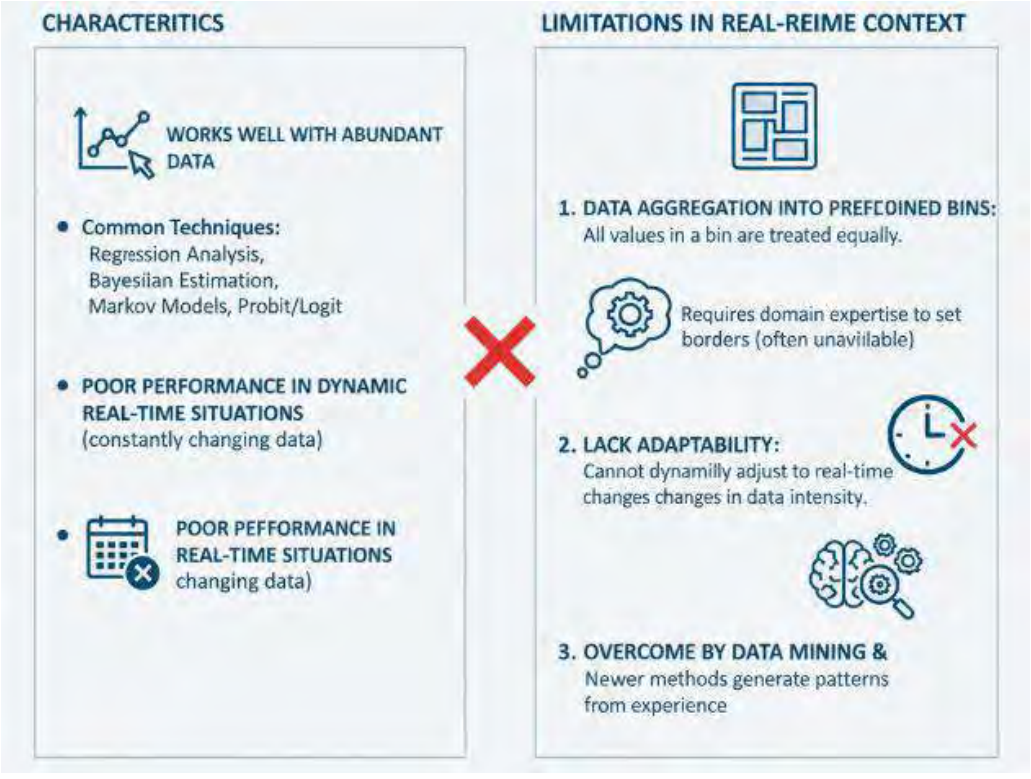

Fig 8 . 4 : Traditional Statistical Methods: Limitations in Real-Time Systems

8.6.3. Hybrid Models

Hybrid Models Machine learning techniques offer a general framework for selecting and combining individual rules for a classification problem. The most commonly applied method for fraud detection is logistic regression, which offers probability and confidence estimates for classification and incorporates voters with different expertise through weighted voting. More recently, credit card companies have increasingly turned to neural networks for fraud detection, capable of reliably detecting fraud when provided with representative data. A crucial challenge in training these systems is the rapid response to fraud patterns before they become irrelevant.

Statistical and machine learning techniques have gradually been augmented with agent-based models. In such agentic systems, various specialized agents receive information about a transaction, whose type depends on the problem. Each agent evaluates the transaction and contributes its assessment to a weighted sum, compiled by a manager agent. This approach has successfully detected fraud in financial transactions and explored it in other domains. Figure 6.7 illustrates fraud detection setups. Specialized biddings for a fraudulent transaction are classified by a voting algorithm, and agents independently assess the likelihood of fraud. Rapid assessment is imperative because, once a fraudulent transaction type is identified, other fraudsters swiftly adapt to evade detection.

8.7. Future Directions

Fraud detection remains an active research area with promising directions. Emerging technologies such as blockchains and quantum computing hold particular promise for advancements in real-time fraud detection. These technologies could significantly boost the performance of existing agentic systems and play a pivotal role in the next generation of intelligent agentic systems.

The application of agentic systems in real-time fraud detection, especially in payment ecosystems, also represents a viable avenue for investigation. Advances in real-time data streaming, processing, and analysis are expected to broaden agentic capabilities and improve their effectiveness in combatting fraud."

8.7.1. Emerging Technologies

The advent of wireless devices, connected vehicles, smart cities, and smart homes requires the collection and processing of massive amounts of data. These new data sets provide numerous possible opportunities and address multiple research questions for problems and challenges in many domains such as climate, healthcare, mobility, etc.

Along with these challenges for many sectors, these domains face issues of fraud, attacks, thefts, terrorism, and so forth. Therefore, it becomes crucial to develop fraud-attack-detection methods and techniques to identify frauds and attacks in real time. The development of fraud-detection systems needs to address new challenges of real-time data-processing capabilities for classification and decision-making. Recent advancements in agentic systems provide the autonomous decision-making capability to classify data elements as fraudulent or normal.

Agentic systems are typically defined as automated approaches that demonstrate autonomous decision-making capable of performing specific tasks with distinct levels of human-like characteristics. The article delves into the capabilities of agentic systems, outlining the requirements of real-time fraud-detection systems and the necessity of real-time data processing for fraud detection. It culminates with a discussion of future research directions that may enable real-time fraud detection through agentic systems.

8.7.2. Potential Research Areas

Research in real-time fraud detection focuses on several critical areas. Algorithms, data processing, and agentic systems have been separately addressed in many academic papers [5,7,9]. The advantages of integrating agentic systems with real-time data processing have been stressed. Future directions may suggest using agentic systems for automatic database record updates.

Real-time data processing has indeed attracted researchers' attention in recent years. Transformed data must be processed in real or near-real time for certain applications. The selected application will then perform an action, provided there is sufficient time. Projects have been designed and implemented to analyze Twitter data in real time. However, real-time data processing introduces some additional problems. Requesting data from the Internet in near-real time can be challenging due to varying Internet speeds and latencies. Unlike historic data, the number of records requested from the Internet at any point in time cannot be precisely controlled; it depends on the Internet speed. Additionally, the time required to analyze different records varies—some may be analyzed within seconds, while others take nearly five minutes. These issues have been acknowledged but not addressed in the experiments.

8.8. Conclusion

The rapidly evolving financial technology landscape underscores the importance of real-time fraud detection and prevention. Implementing agentic systems enables organizations to counter growing volumes and complexities of fraudulent transactions

effectively. Agentic systems can establish highly automated workflows—ranging from using their environment to trigger complex decision-making against cyber-related fraud to updating and maintaining their influence upon the transactional environment via state-of-the-art cyber defense operations—thereby decreasing the operational gap faced by traditional methods. The dynamic combination of agentic systems and real-time processing further enhances the ability to detect and prevent fraudulent transactions promptly, preventing exploitation, financial losses, and weakening the economy.

Fig 8 . 5 : Fraud Incidents Over Time: Agentic vs Traditional

In summary, the growing threats of fraud, identity theft, and theft of personally identifiable information, coupled with the increasing sophistication of attack methods and exploitation of zero-day vulnerabilities, accentuate the criticality of real-time fraud detection. Organizations capable of executing real-time monitoring benefit from a narrowed operational gap and increased capability to detect and halt fraudulent transactions rapidly, stemming potential financial and data-related damages.

8.8.1. Summary of Findings and Implications

Agentic systems represent a distinctive category of information systems that demonstrate a degree of autonomy through the employment of artificial intelligence, automation, and related technologies, enabling what might be characterized as self-directed behaviour. These systems possess either considerable authority or recognised agency and thus the capacity to perform specific tasks. The rise of agentic systems has prompted several business case studies and in-depth investigations into their behaviour, encompassing various physical, temporal, role-related, and socio-cultural characteristics, and examining special cases such as chat bots and robotic process automation. Given their

broad applicability, agentic systems have been used to address critical risk domains, including fraud detection.

Fraud detection remains a high priority for many organisations. A recent trend involves advancing beyond post-event fraud processing towards the real-time detection of potentially fraudulent exploits. The continuous availability of power, connectivity, and data—characteristic of modern digital infrastructures—underpins this shift. Techniques to collect and stream transactional data, enabled by developments in Internet of Things technology, allow the maturing of real-time fraud detection methods and algorithms. Despite extensive research on fraud detection and the underlying algorithms, the real-time aspect—especially when agentic systems are employed—has yet to be thoroughly explored. The intersection between the agentic nature of a system and its capacity to readily respond to streaming transactional data constitutes a fertile area for further study.

References

[1] Diez O. (2014). Resilience of Cloud Computing in Critical Systems. Quality and Reliability Engineering International, 30(3), 397–412.

[2] Meda, R. (2025). Integrated Sales Performance Management Platforms: Leveraging AI for Quota Allocation, Demand Forecasting, and Zone-Based Sales Optimization. Advances in Consumer Research, 2(4).

[3] Tärneberg W, Skarin P, Årzén K-E, Kihl M. (2023). Resilient Cloud Control System: Realizing Resilient Cloud-Based Optimal Control for Cyber-Physical Systems. arXiv, 2304.00857.

[4] Kalisetty, S., & Inala, R. (2025). Designing Scalable Data Product Architectures With Agentic AI And ML: A Cross-Industry Study Of Cloud-Enabled Intelligence In Supply Chain, Insurance, Retail, Manufacturing, And Financial Services. Metallurgical and Materials Engineering, 86-98.

[5] Wijaya S, Ramadhan A, Andhika A. (2023). Cloud-Based Control Systems: A Systematic Literature Review. International Journal of Reconfigurable and Embedded Systems, 12(1), 135–148.

[6] Kalisetty, S., Preethish Nanan, B., Annapareddy, V. N., Gadi, A. L., & Kommaragiri, V. B. (2022). Emerging Technologies in Smart Computing, Sustainable Energy, and Next-Generation Mobility: Enhancing Digital Infrastructure, Secure Networks, and Intelligent Manufacturing. Venkata Narasareddy and Gadi, Anil Lokesh and Kommaragiri, Venkata Bhardwaj, Emerging Technologies in Smart Computing, Sustainable Energy, and Next-Generation Mobility: Enhancing Digital Infrastructure, Secure Networks, and Intelligent Manufacturing (December 15, 2022).

[7] Shirazi N-u-h, Simpson S, Oechsner S, Mauthe A, Hutchison D. (2015). A Framework for Resilience Management in the Cloud. Elektrotechnik und Informationstechnik, 132, 122–132.

[8] Gadi, A. L. (2020). Evaluating Cloud Adoption Models in Automotive Manufacturing and Global Distribution Networks. Global Research Development (GRD) ISSN: 2455-5703, 5(12), 171-190.

[9] Konaganti S.D.P. (2023). Intelligent Resilience in Multi-Cloud Systems. International Journal of Scientific Research in Science and Technology, 8(6), 431–440.
[10] AI-Enabled Predictive Modeling for Flood and Mobile Home Insurance Claims Management. (2025). MSW Management Journal, 34(2), 1295-1316.