



Autonomous Fraud Defense for Cloud-Native Insurers

1st Phanish , 2nd Hiroshi Tanaka

Phanish Lakkarasu, Senior Site Reliability Engineer , phanishlakkarasu@gmail.com,
Hiroshi Tanaka , Research Assistant

Abstract

Fraud remains a critical challenge for the insurance sector, which increasingly relies on Data-Driven Decision systems for its identification and prevention. Cloud and Data Ecosystems now enable the use of expanded sensor networks, supporting smarter and more responsive insurance products. An evolving Cloud-Native Architecture built on Fraud Detection-as-a-Service allows fraud detection Agentic services to be rapidly deployed and decommissioned as needed. However, two critical dimensions—the scaling of Agentic properties and the incorporation of settlement Risk from False Positives—are often neglected or addressed too late in Cloud-Native Fraud Detection-as-a-Service implementations, even though Agentic properties are essential to scaling these services effectively in the cloud.

To address this gap, Causal Advisory-Multi-Method Evaluative Agentic Fraud Detection offers a framework for incorporating Agentic properties into Cloud-Native Fraud Detection-as-a-Service, applying a risk-causal perspective to their evaluation and modeling. Complementing this, a Cloud-Native Framework for Agentic Data Governance is proposed to ensure Data Quality, Privacy, and Regulatory Compliance across Fraud Detection-as-a-Service operations, while also enabling ongoing monitoring of settlement Risk from False Positives. The resulting architecture follows a Data-Pipeline-driven deployment pattern, built on a Microservices base with a Serverless Control Plane.

Keywords: Insurance Fraud Detection, Fraud Detection-as-a-Service (FDaaS), Cloud-Native Fraud Architecture, Agentic AI Systems, Agent-Based Fraud Detection, Causal Fraud Modeling, False Positive Risk Management, Data-Driven Decision Systems, Real-Time Fraud Analytics, Microservices-Based Architecture, Serverless Control Plane, Data Pipeline Architectures, Cloud Data Ecosystems, Data Governance Frameworks, AI-Driven Risk Detection, Privacy-Preserving Fraud Analytics, Regulatory Compliance in Insurance, Scalable Fraud Detection Systems, Intelligent Risk Advisory Systems, Autonomous Fraud Detection Agents.



1. Introduction

Agentic AI-Driven Fraud Detection and Autonomous Decision Systems in Cloud-Native Insurance Platforms: present a concise, evidence-based, policy- and practice-oriented synthesis with formal structure and rigorous references.

Insurance fraud threatens all stakeholders. Insurers suffer financial losses and damaged reputations; innocent policyholders bear the burden of contended premiums and inadequate policy payouts; and external parties become unintended victims. Rule-based fraud detectors can be readily defeated. Anomaly-based detectors exhibit high false-positive rates, while machine learning-driven detectors require substantial training data and achieve limited generalizability. Agents equipped with artificial intelligence and operated at higher levels of autonomy present an alternative that, hitherto, has not been rigorously studied.

Are agentic AI systems a viable solution to the lifecycle management of detection within cloud-native insurance platforms? Addressing this question requires establishing a formal foundation for agentic AI in the insurance domain and determining the implications of their use for fraud detection systems. A conceptualization of agentic AI, mapping agency characteristics, the decision-making process controlled by agents, and the insurance domain forms the basis for studying agentic fraud detection. The approach transcends information technology and system lifecycle boundaries; in terms of the latter, it emphasizes data ecosystems.

2. Theoretical Foundations of Agentic AI in Insurance

A conceptual model links agentic intelligence, autonomous decision systems, and business outcomes in insurance contexts.

The relationship between agentic AI and business outcomes within the insurance sector has been a subject of considerable theoretical debate in recent years. The prevalent view posits that agentic systems should primarily assess operational risk, allowing agents to respond with externally controlled decisions, supplemented by IT-enabled automation. An emerging counter-narrative proposes a different perspective: the internal decision system of AI fraud detection solutions within cloud-native insurance platforms should be empowered to autonomously detect and mitigate fraud using a full agentic intelligence capability set, thereby addressing operational risk. The core contention revolves around the definition of agentic AI and its actual complementarity with respect to business outcomes. Rather than being purely aligned with risk



reduction, the internal agentic intelligence of cloud-native insurance platforms should be viewed as potentially enabling enhanced operational and market performance.

This broader definition directly influences the governance of internal AI decision-making capability. Where the previously established view requires fraud detection with full autonomy solely for minimizing operational risk, the more inclusive perspective on operational control implies that the external alignment or safety of the fraud detection system is a necessary and sufficient condition for effective operational governance. The expanded view permits a related broadening of the agentic intelligence concept and its governance requirements, not just in the context of basic model performance but with respect to any operational dimension or NPM-associated performance measure. Finally, a conceptual model that maps the relationship between agentic intelligence, autonomous decision systems, and business outcomes in insurance contexts is delineated.

2.1. Agentic Intelligence and Autonomous Decision-Making

Agentic AI is a specific form of artificial intelligence that relies on agentic decision-making processes in which the system is cognitively capable and responsible for key decisions. Three main characterizations of agentic decision-making can be considered: levels or associated aspects of cognitive and operational autonomy; decision loops and execution control processes that can be defined for decision execution; and the linking of risk associated with these systems, including risk prioritization in subsequent decisions. Although the concepts addressed can apply to other domains, these aspects of agentic AI are formalized with a specific focus on their application in the context of insurance, reflecting the growing interest in these types of solutions in this domain, particularly for processes related to threat detection.

Agentic AI in insurance can be defined as those cognitive AI systems where human input is no longer needed or needed only at the top levels of supervision or governance of the system. In this case, the levels of autonomy of the cognitive AI system installed in the insurance company are appropriate for the company to accept the risks associated with the decisions made; systems are deployed in controlled environments that allow for user privacy, in accordance with applicable legislation, industry norms, and company policy; and the adequacy to the business and freedom from undesirable discrimination are analyzed and guarantees established through suitable methods. It has also been pointed out that the risk involved in agentic fraud detection systems with the capacity to operate close to real time and to perform an independent decision



loop is also part of the system and, as such, should also be controlled. The development of empirical knowledge and general guidance in this area can therefore be useful to those constructing and deploying such systems.




Fig 1: AI agents for fraud detection

2.2. Fraud Detection Paradigms in Insurance

The use of fraud detection systems has been on the rise in the insurance industry as the volume of fraud increases. While these detection systems are designed to help flag fraudulent claims for further investigation, they are not without their disadvantages. As more insurance companies use fraud detection systems, fraudsters are taking advantage of those weaknesses; overcoming these weaknesses is essential. In insurance fraud detection, there are four main paradigms for detection: rule-based detection, anomaly detection, machine learning-based detection, and hybrid detection. Rule-based detection systems are the oldest forms of fraud detection systems. They validate claims against a preset list of rules and regulations, flagging any claims that do not comply with the internal regulations defined by the insurance organization. In anomaly detection, the historical data of legitimate claims are considered, and an anomaly score



is assigned to each new claim based on this data. If the anomaly score is above a defined threshold, the claim is flagged for investigation. Machine learning-based detection is a newer, more effective approach where historical fraud data are used to build a model that is able to identify fraudulent claims. Hybrid approaches combine machine learning and rule induction systems.

Although fraud detection methods can vary, they should be evaluated based on the same criteria. Detection accuracy, the percentage of flagged fraudulent claims that are indeed fraudulent, suffers from the issue of class imbalance in the fraud detection domain. Latency is also important in fraud detection; lag in flagging fraud is costly, considering the goal of reducing financial losses. Missing fraud detection can cause serious losses for insurance companies. If a detection system is constantly flagging false positives, it will lead to wasted resources reviewing these claims. Consequently, the false-positive rate of a fraud detection system should always be kept low, as a high false-positive rate will also lead to increased costs of claim handling and can cause customer dissatisfaction and loss of customers.

Equation 1. Fraud probability score

Start with claim features:

$$X = (x_1, x_2, x_3, \dots, x_n)$$

Assign model weights:

$$W = (w_1, w_2, w_3, \dots, w_n)$$

Compute raw fraud risk:

$$z = w_1x_1 + w_2x_2 + \dots + w_nx_n + b$$

Compact form:

$$z = W^T X + b$$

Convert to probability using sigmoid:

$$P(\text{Fraud}|X) = \frac{1}{1 + e^{-z}}$$



Final:

$$P(Fraud|X) = \frac{1}{1 + e^{-(W^T X + b)}}$$

3. Methodology

A multi-method research design supports the analysis. The primary case study—an industrial cloud-native insurance platform—provides contextual insights and examples. Technical models are grounded in secondary data and draw upon practitioner experience. Formal risk and threat analyses offer theoretical validation. Agentic AI’s application to fraud detection is explored through schema-definition approaches and proof-of-concept development.

The Research Framework for Data Governance in Cloud-Based Systems specifies key elements of data governance in a Cloud services context—namely, stewardship, data lineage, access controls, data retention, and regulatory compliance. Data stewardship assigns accountabilities for data quality and privacy compliance, while data lineage tracks data provenance across systems and regions. Access controls restrict data exposures to mitigate privacy risk in normal operation and during data breaches, analyses, and development. Data-retention policies specify retention periods for normal operation and the duration of exception removals. Compliance pathways address liabilities and remedial responsibilities relating to GDPR and AI Act non-compliance and breach notification requirements under GDPR Article 33.

Table 1: Core Concepts and Terminology in Agentic Fraud Detection

Term	Meaning	Role in System
Agentic AI	Autonomous decision-making AI	Performs fraud detection without human intervention
FDaaS (Fraud Detection-as-a-Service)	Cloud-based fraud detection platform	Enables scalable deployment
False Positive Risk	Legitimate claims flagged as fraud	Impacts cost and customer trust
Data Governance	Data control (lineage, access, retention)	Ensures compliance and privacy



Term	Meaning	Role in System
Microservices Architecture	Modular service-based design	Supports scalability and flexibility
Serverless Control Plane	Event-driven execution model	Enables automatic scaling
Drift Monitoring	Detects changes in data/model	Maintains model accuracy
Auditability	Traceability of decisions	Ensures transparency and compliance

3.1. Framework for Data Governance in Cloud-Based Systems


A framework for data governance in cloud-native systems defines responsibilities across platforms, consumers, and regulators. Data stewardship tracks lineage, enables policy enforcement, and drives compliance. Stewardship records document ownership, quality controls, and risk mitigation. Comprehensive data lineage tracks information from provenance through delivery to expiry, ensuring accessibility to all relevant stakeholders. Data access controls govern visibility and modification rights, with separate pathways for test and production datasets. Explicit retention policies define data lifecycle stages and disposal criteria.

Cloud-based environments support rapid innovation and dynamic business models but create new governance challenges. Non-hierarchical ownership models complicate data-management responsibilities, while external storage and processing introduce additional operational risk. The use of data from diverse external sources creates compliance concerns related to third-party agreements and privacy legislation. Data provenance dictates appropriate protective measures against misuse and ensures compliance with vertical- and horizontal-privacy requirements. Residual risk inherent in permissionless access and external sharing requires mitigating controls.

Available data lineage may also be incomplete or poorly documented, diminishing usefulness. Information from external sources may introduce hidden replacement costs that are not reflected in price-sharing models. Information-monopoly situations impose additional risks. The absence of formal data-access controls may cause test or other non-production data to flow into production use, resulting in unintended consequences. Specified controls ensure ongoing



data governance and compliance and provide necessary assurance to external stakeholders and regulators.



4. Objective of the Study

The main objective is to present an investigation of agentic AI in the context of fraud detection for cloud-native insurance platforms. Such environments serve diverse objectives that give rise to distinct detection models, each explicitly optimized for a specific aspect of the task. While each approach handles a fragment of the threat landscape, they can be cascaded into a coherent fraud scoring system that provides support to the business rules. Academic and practical literature on agentic AI is vast, but a formal synthesis that brings together the key aspects from different works serves to clarify the concept and make it amenable to systematic investigation and deployment.

Agentic AI facilitates a new class of nondeterministic decision systems capable of formulating hypotheses in uncertain environments through data-driven learning. Such systems



generate hypotheses that explicate the feature interdependencies affecting the outcome and assign an implausibility score to each observation. Consequently, agentic AI enables decision systems to operate under nondeterministic environments that, based on training, test, and validation data, inform the boundary between plausible and implausible behavior. Although any decision system can benefit from explanation-assisted transparency (for instance, to enable human–machine collaboration), agentic AI is particularly relevant for distributed decisions that lack human input or oversight. The introduction of the concept in the insurance domain enables the creation of more effective solutions to the monitoring, detection, prevention, and investigation of fraud.




Fig 2: Key Capabilities Of Agentic AI

4.1. Aim and Scope of the Research

The research seeks to derive a framework for agentic fraud detection systems in cloud-native insurance platforms and investigate specific analytical, monitoring, and validation aspects. Cloud-native insurance platforms embrace new technologies to offer substantial user benefits, but they also expose their ecosystems to novel fraud risks. These adversities can be met by



scaling, orchestrating, or making fraud detection systems agentic. Agentic AI embraces higher-level generalization and reduction of supervisory effort through increased autonomy. A common cloud-native architecture pattern is based on the microservices or serverless paradigm. For agentic fraud detection systems, cloud-native patterns comprise agentic model development, testing, and operational lifetime management.

The research framework-oriented structure ensures coherent coverage of technological, methodological, and validation aspects. New methodologies for supporting a specific aspect of the analytical process are proposed. The automation of dynamic monitoring of changing risks and the evolution of effective countermeasures are presented, together with patterns for rigorous validation of ML models, including chaotic and agentic AI systems. The framework, although specific to fraud analysis, encompasses categories relevant to agentic systems generally.

5. Research Summary

Core findings and contributions are presented in a framework that addresses stakeholder roles, required triadic relationships for data ecosystem governance, and salient interoperability requirements for agentic AI applications. Key data management issues—ecosystems supporting immunity to privacy risks and regulatory breaches; accurate secretive state estimation and mitigation of associated dangers; and seamless consensus establishment amid transitive data flow—are addressed.

To enhance an insurance business's advantage, core operational systems must deliver the utmost accuracy, speed, and efficiency. Demand is thus emerging for state estimation and real-time scoring systems capable of detecting operational threats as they emerge, enabling deployment of mitigation resources pre-emptively and automatically. Notwithstanding the clear business benefits, the underlying data ecosystems must be managed effectively to avoid infringing third-party immaterial rights. An agentic cloud-native personal-lines insurance-fraud-detection deployment illustrates how harmony can be achieved, producing immunity to all known privacy and regulatory breaches by establishing clear paths for managing data consent, transitive flow, and recipient control. Enabling agentic AI-driven systems requires clearly establishing participants' data alignment and the necessary healthy triadic discussions.

Equation 2. Fraud decision rule

Let fraud threshold be:



τ

Decision:

$$\hat{y} = \begin{cases} 1, & P(\text{Fraud}|X) \geq \tau \\ 0, & P(\text{Fraud}|X) < \tau \end{cases}$$

Where:

$$1 = \text{Fraud}, \quad 0 = \text{Non-Fraud}$$

5.1. Data Ecosystems: Navigating Privacy and Compliance Challenges

Rising regulatory scrutiny on cross-border data transfers, privacy protection, and algorithmic bias creates significant challenges for insurance data ecosystems. Complex data provenance, as multiple parties contribute data from different jurisdictions, increases the likelihood of noncompliance with data governance frameworks and privacy legislation. In addition to information belonging to external entities, such as policyholders, the insurance industry relies on third-party data for marketing, underwriting, pricing, claims detection, and fraud investigation. The potential harm from an unintended disclosure of sensitive data, such as health and financial information, is considerably high. Therefore, processing sensitive data requires robust protection mechanisms and compliance with data protection legislation.

Data ecosystems also have an inherent privacy risk: algorithms of sufficient power capable of generating high-accuracy predictions can inadvertently learn and retain sensitive information about members of a protected sub-group, and, even when the training data is anonymized or rendered insensitive, the preservation and disclosure of any sensitive information with respect to these individuals could lead to discrimination. Moreover, the definition and implementation of a consent mechanism governing the use of members' data within the data ecosystem add to the difficulty of meeting regulatory requirements. With multiple entities providing and utilizing cross-border data origins for a common purpose, collecting user consent for the processing of each individual cross-border data transfer of data sources is highly impractical and increasingly being perceived as unreasonable by users. Accordingly, the inclusion of external data—especially sensitive external data—should be limited to the absolute minimum, minimizing cross-border transfers of sensitive external data. Furthermore, data



minimization and anonymization strategies should be adopted when integrating such sensitive data from unknown sources.

6. Data Ecosystems and Privacy Considerations

Privacy and compliance concerns permeate many aspects of cloud computing architectures. The distribution of infrastructure across multiple geographies poses challenges for national privacy legislation, as data lawmakers have little control of the destination of their citizens' data. Large data migrations over many jurisdictions raise the risk of noncompliance, and thus of heavy legal fines. As the size of the data lifecycle increases so does the potential for improper publication of sensitive information. Moreover, many business models leverage third-party data aggregation and searching capabilities.

Privacy, security, and compliance consider privacy at scale from national privacy laws to proper information leak controls; ensure that potential and perceived dangerous information cannot be made available by any means, even indirectly; protect sensitive data with confidentiality mechanisms and governance policies that provide appropriate access control; and ensure that deep learning models do not memorize sensitive data, are robust against adversarial attacks, and mitigate all potential points of information leak.

Table 2: Fraud Detection Paradigms Comparison

Method	Accuracy Level	False Positive Rate	Latency	Adaptability
Rule-Based Detection	Low	High	Low	Very Low
Anomaly Detection	Medium	Medium-High	Medium	Medium
Machine Learning	High	Medium	Medium	High
Hybrid Models	Very High	Low	Medium-High	High
Agentic AI Systems	Very High	Very Low	Low	Very High

6.1. Data Governance and Quality in Cloud-Native Environments

A comprehensive data governance framework tailored to cloud-native insurance ecosystems fosters data quality, structured metadata, efficient data quality metrics and thresholds, and delineated stewardship roles. The framework addresses data ownership and



stewardship, lineage management, information protection and access control, data retention and deletion, and cross-border compliance.

Cloud-native computing changes how organizations build and procure technology and underpin them. Multi-cloud architectures, with mostly private, a hybrid set-up, or interconnected cloud vendors, offer greater agility and cost savings to sensitive industries such as banks contrasted with commoditized service sectors leveraging public cloud.

Yet, cloud technology adds pressure on the third-party risk management and mitigative strategy of companies, driving around 25% of expenditure under consideration into compliance organisations focusing solely on third-party risk.

Appropriate safeguards must cover all services and deployments for preserving data.

The insurance industry must encourage changing the attitude of sharing data. The old gossip that “data is the new gold” is not necessarily true; instead, “data is the new oil” is all but accurate.

An understanding exists that data producers, recipients and intermediaries, or owners, authorities and regulators, must share the data for greater benefits, and not just hiding and protecting it within silos.




Fig 3: The 4-Step AI Workflow For Fraud Prevention

6.2. Privacy, Security, and Compliance in Fraud Detection

The security of data, such as encryption schemes for protection during transmissions, storage, and processing at cloud providers, must be assured to address growing concerns over thefts, vulnerabilities, and misuse affecting customers' goodwill and trust. Access governance requires strict authorization to prevent unauthorized access, revealing privacy-sensitive information. Appropriate modelling using the principle of least privilege helps to establish stringent access governance policies. Security requirements are paramount when agentic systems are exposed to external data sources and destination systems. These external data sources and destination systems are potential sources of new classes of attacks, and so adapting the agentic system's behaviour to avoid or recover from these attacks is crucial. Threat modelling explains potential security threats and validates the efficacy of the agentic system in repelling these attacks. Privacy experts recommend audits to assess whether privacy risks governed by privacy policies are acceptable to all stakeholders. Auditable models track users' trust relationships. Using a distinguished, trusted agent which monitors movie rentals, peer-to-peer transactions



among users are subject to embargo and are not viewed by landlords, but audited later. The principle of least surprise helps structuring a useful audit trail so that users perceive any audit as a minor and expected breach of privacy.

7. Methodologies for Agentic Fraud Detection

Approaches to modeling, monitoring, and evaluation derive from the theoretical foundations. Causal inference concerns about biases shaping predictions have led to checks of the conditions for drawing causal conclusions beyond mere correlations. Given the conscious and subconscious nature of human decisions and judgments, there are also considerations addressing the difficulty of producing understandable post hoc explanations across black-box agentic deep-learning systems through humanly interpretable rule systems or simplified models simulating the complicated processes.

Discriminative model architectures that build a real-time score on an insurance-related transaction raise specific aspects of the modeling performance. It is infeasible to have latency-sensitive deployment as consolidated tables from the data warehouses that underpin SSAS cubes offer latency in the order of hours. The data pipeline built to maintain the daily data drift technique demands seven test populations to record each day the model training with six test populations, and the last point is for validation. When a high number of features is tested, the daily drift can bloat the process with building models based on synthetic populations from cross-validation, and the common solution to counter this problem is through online learning or decreasing the number of archetypes or features.

Equation 3. Accuracy

Correct predictions are:

$$TP + TN$$

All predictions are:

$$TP + TN + FP + FN$$

Therefore:



$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$


7.1. Model Architectures for Real-Time Fraud Scoring

The core requirements of a model architecture for real-time fraud flagging include a low latency from the submission of relevant feature data to the return of a score or class label, high reliability in scoring correctly classified individuals or transactions as ‘not fraud’, and the ability for the predicted scores to correlate closely with the estimated fraud risk according to the trained model. Recent papers have compared the performance of some of the more popular models used for fraud scoring and shed some light on these requirements; comparisons across other types of time-limited fraud prediction in different domains also reveal relevant information. The aspects that could contribute to lower time-to-score performance include the use of relatively shallow architectures (only a few layers), the investigating of more constrained architectures that resort to ensemble learning only at the last layer, and the incorporation of non-standard computational resources—specifically, by resorting to GPUs rather than CPUs.

The monitoring of live fraud model performance can happen in various ways: by computing streaming model quality measures directly on the data stream in production; through the comparison of a designated test set with a training set, plus the appropriate drift detection tests; and via a dedicated departitioning strategy. The latter approach allows for detecting changes in the separation of classes detected by different fraud models and has been applied successfully in financial settings. Deploying a separate watch-list fraud model acting as a first gate assessing a transaction or customer for fraud risk, in combination with successive models processing the other transactions, would help to increase throughput further (albeit not in an efficient manner) without compromising quality.



Graph 2: Cloud-Native Agentic Fraud Architecture Components



7.2. Causal Inference and Explainability in Autonomous Systems

The absence of human sensibilities in agentic AI can complicate the evaluation of fairness, ethical conduct, and accountability in autonomous systems. Causal graph representations can address these issues, but they depend on a priori knowledge. In the absence of adequate causal knowledge, explanatory methods that highlight anomalies in the decision processes can help. Counterfactual explanations present information on the influence of input variables on a particular outcome. Textual presentations provide a natural and agile approach to counterfactual explanations.

In contrast to supervised learning techniques, which formulate the decision rules of the models, analogy-based reasoning only confirms that the absence of the process of comparison and contrasting, which is inherently human. In a scenario in which the two concepts are isolated, analogy-based reasoning with a supervised learning technique derives similar conclusions to those arrived at through supervised learning. The human ability to express reasons for a particular decision cycle requires a deeper understanding of the origins of those decisions. It is



this ability to inform the relevant reasons for a specified decision, as in expediting a task or the omission of a task, that is still lacking in existing autonomous agentic systems. Biological agents convey usefully informative messages about decision making, thus explicating the decision process and bringing agents one step closer to human sensibility.

Table 3: Evaluation Metrics for Fraud Detection Systems

Metric Name	Formula	Purpose
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Overall correctness
Precision	$TP / (TP + FP)$	Correct fraud predictions
Recall (TPR)	$TP / (TP + FN)$	Detection of actual fraud
False Positive Rate	$FP / (FP + TN)$	Wrong fraud flags
F1 Score	$2 \times (Precision \times Recall) / (Precision + Recall)$	Balance metric
Latency	Total processing time	Speed of detection
Privacy Risk	Data exposure probability	Compliance measure
Auditability	Traceability score	Trust and governance

8. Trust, Ethics, and Governance of Autonomous Insurance Systems

The evolution of autonomous insurance systems toward true agency triggers doubt and resistance, partly due to how decisions converge with sensitive societal and personal domains of life. The discussion centers on accountability for actions governed by these systems and extends to external influences—whether from intentional, malicious alteration or inadvertent manipulation through bias in the supporting data. Similar considerations arise regarding support and surveillance of advanced autonomous systems controlling dangerous civil-energy applications. However, in the case of insurance, human involvement traditionally looms large.

Recognizing policyholder behavior and anticipating customer needs had long been central. Empathy for policyholders was deemed essential in direct claims settlements; professionalism, trustworthiness, and duty were invoked in regard to brokers. The question, different in emphasis though not in substance, is therefore whether a human being at the helm must always be the focus—and whether external control and guidance grounded in agency



theory, transparency, and accountability can offer enough assurance, in which case a shift toward a complete reliance on techniques, models, data, and automated processes may genuinely usher in the holy grail of demands: speed, efficiency, fraud elimination, risk-based pricing, and hugely lower premiums.

8.1. Accountability and Liability in Agentic Systems

Responsibility for insured interests lies primarily with people. While growing risks and their high-cost consequences must be borne by the entire community, the party that appears to be liable should bear that risk. Agentic systems do not assume such responsibilities; their role is to assist humans in making better decisions. This means assisting humans in generating real control over those agentic systems engaged in becoming more autonomous. Transparency and auditability are key here; they help understand how these agentic systems operate, thus ensuring accepted standards of control even if fielded in a regulated condition. In this task, human supervision is critical; it cannot be neglected, as its absence can be the cause of potential undesired outcomes.

Via the latest technological breakthroughs emerging from the AI domain, supervision can become even more effective, through decision-centered mechanisms that support about explaining into the AI domain and generating *fiabilité explicative*; and such opportunity can be fully leveraged through the integration of cross-disciplinary knowledge collected in both soft and hard-scientific fields offering interesting breakthroughs to understand the human decision chain better and its relationship with the agentic systems supporting the decision. Still, the autonomous agent must create reliable decisions with respect to the agents compared, even if these decisions are, at most, “supportive”. This creates the need to clarify clearly the boundaries of upcoming liabilities basically originating from human decision-makers: first when onboarding an agentic system of any kind, and second when jump-starting its functioning.



Fig 4: Agentic AI Vs. Generative AI: Redefining Fraud Detection

8.2. Fairness, Bias Mitigation, and Inclusion

Imbalances in decision systems can originate, be propagated, or be exacerbated through AI algorithms. The sources of these imbalances include difficulties in obtaining sufficiently diverse data, unintentional biases in data, such as when individuals from protected groups have fewer examples of a given risk behavior, and deficiencies in the development process. In areas such as recruitment or lending, where protected groups are represented in sufficient numbers, bias can be controlled by simply analyzing results stratified according to the sensitive attributes. However, this option is generally not available for fraud detection because individuals from protected groups usually represent only a small percentage of the overall data set. The most commonly applied techniques to mitigate bias are preprocessing, in-processing, or postprocessing strategies that adapt the training or prediction process. If potential bias has been recognized, documented, and analyzed in detail, then bias mitigation techniques can systematically be applied.

Although fairness is normally described in terms of a single sensitive attribute, an increasingly recognized approach assesses bias against a base risk behavior defined by



combining multiple attributes. A similar multi-attribute technique could apply to a fraud-scoring model. For example, in addition to gender and ethnicity, the set of sensitive attributes might comprise aspects such as location or vehicle ownership. Bias analysis can also extend to not only the predicted fraud class but also the risk score. Standard evaluation metrics for predictive bias focus on a subgroup G of a protected group and measure the comparison with the rest of the group (denoted $\bar{G}$). The bias analysis may be non-binary, whereby each subgroup G is compared with the rest of the population.

9. Cloud-Native Deployment Patterns for Agentic Fraud Detection

Agentic AI is a central theme of the agentic fraud detection paradigm. Therefore, it needs to be considered not only from a logical standpoint, as done in Methodologies for Agentic Fraud Detection and Model Architectures for Real-Time Fraud Scoring, where the focus was placed on defining the modeling techniques and their evaluation, but also with respect to deployment architecture. This section describes the cloud-native deployment patterns that allow agentic fraud detection systems to scale according to the availability of data in a cost-efficient manner.

Although the flexibility and cost efficiency of deployment is important for other detection approaches as well, agentic systems need to present additional requirements that arise from the presence of decision loops. The effects of these decisions must be limited; hence, monitoring and auditing aspects become essential.

Agentic detection systems consist of any detection technology, deployed in a full AI-as-a-Service mode, that models the agentic properties of fraud. These schemes are layered compositions of fraud detectors that, due to their nature of agentic detection, break the cyclic dependency on post hoc mitigation measures of traditional approaches. They guide the decision agents (automated or human) during intervention with predictive fraud scores. The scoring must be done in real-time, meaning that the prediction should occur at the time of user interaction, transaction, or information consumption. Real-time prediction usually comes with latency constraints that become stricter when the number of concurrent requests increases significantly.

Equation 4. Precision

Fraud cases flagged by system:

$$TP + FP$$



Correct fraud flags:

$$TP$$

Therefore:

$$Precision = \frac{TP}{TP + FP}$$

9.1. Microservices and Serverless Architectures

Cloud-native fraud detection systems can leverage the full range of cloud service models, but service-oriented components often yield the greatest flexibility, scalability, and resilience. A microservices architecture employs lightweight, granular services with focused responsibilities, enabling rapid agile development; it enables horizontal scaling to manage variable workloads; and it minimizes the risk of systemic failure. Serverless architectures abstract away server management entirely, automatically provisioning backend resources in response to demand, creating a strongly consumption-based cost model, and freeing development teams to concentrate on writing code without operational concerns. However, the invocation latency of serverless components must be considered when architecting a complete deployment.

Individual services must be capable of connecting to one another securely and seamlessly. Embedding the necessary parameters—such as keys—directly in the service is not advisable as it violates the principle of least privilege; far better to provision security features such as message brokers, orchestrators, and secret stores. Microservices are also typically deployed in containers, and container orchestration tools such as Kubernetes help deployment, scaling, and operation of such containerized applications. Serverless components generally have an event-driven focus: an event occurs, code is executed to address the event, resources are released. Such architectures can require a service fulfillment back-end, with cloud functions, but processes that follow a workflow can either be managed by an orchestrator such as Apache Airflow, which links the services in their execution order, or be event-triggered.

9.2. Data Pipelines, streaming, and Orchestration

Data travels across multiple systems in the insurance ecosystem. The relevant entities range from external sources, such as social media and postal services, to business partners,



including intermediaries and reinsurers, as well as internal operational systems, such as customer relationships and policies. Data need to be ingested, transformed, and made reachable for processing and storage. This processing often entails enhanced risk scoring, but it may also apply to other situations, such as gathering information for decisions based on supervised classification or causal inference. Most processing is not time-critical, but some analysis generates events that subsequently trigger real-time actions.

The data flows can be seen as pipelines integrating various processes. These pipelines may involve simple extract–transform–load processes or orchestrated execution of complicated business processes involving multiple tasks. Even if not strictly time-critical, the end-to-end latency must still remain acceptable and, therefore, a sufficiently small degree of parallelism is highly desirable. For cases requiring dedicated low latency in response to specific events, the data pipelines must be complemented with streaming, from which out-of-band actions can be triggered. Data must naturally go through these flows multiple times, whether for monitoring, updating scoring tables, or validating activity, and reusable orchestration development and execution functions are therefore required.

10. Evaluation, Validation, and Compliance

Standards for measuring effectiveness, safety, and efficiency comprise detection accuracy, false-positive rates, latency, privacy risk, acceptability across population groups, and safety margin. Acceptance by a recognized standard-setting authority, completion of external auditing, and validation against well-defined testing sets assure all stakeholders that agentic capabilities behave as expected. The architecture of agentic fraud detection systems entails the real-time storage of both internal workings and external decisions in a log; these archives enable after-the-fact explanation and facilitate independent review. Construction and use of a private set of test cases, sufficiently distinct from training samples, assures that validation generates outcomes that maintain internal quality without compromising external fairness.

Cloud-native technologies and practices, when applied to agentic fraud detection systems, facilitate scalability, resilience, and cost-effectiveness while helping repair any control failures. By partitioning the detection logic into microservices and attaching serverless components, organizations can scale selected parts of the processing independently, concentrating computing and storage resources on those processes that require them while relying on the essential spare capacity of the cloud service to deliver cost-effective operations. Event-oriented processing is an



inherent feature of cloud-native architectures; orchestrating decision pipelines based on event-driven design patterns ensures accurate sequence and timing while reflecting the real-world environment of the system.




Fig 5: Adoption Strategy For Agentic AI In Insurance

10.1. Metrics for Effectiveness, Efficiency, and Safety

A multi-dimensional approach to validation is required. Effectiveness is assessed using commonly used measures for detection performance, including accuracy, true-positive rate, true-negative rate, false-positive rate, and F1 score. Support and precision are also critical in an industry where huge resources and costs may be expended to adjudicate and prosecute identified criminals. Indeed, considering the futile and expensive investigations of legitimate customers, it is ideal to have a very low rate of false positives, even at the expense of having the system dismiss some frauds. Latency, and thus efficiency, is also key. Many frauds are committed by teams of people moving quickly to exploit opportunities and using cheap stolen Identifiable Personal Information (IPI) to commit their attacks. Hence, speed is of the essence in detecting fraudulent transactions. Finally, in many businesses it is important to ensure that the actions taken by the fraud detection system do not cause real harm to customers, so an estimate of safety is paramount.

A large section of the industry is currently moving toward achieving privacy-enhancing technologies (PETs) that allow organizations to share information and data analysis results



without revealing sensitive data samples to each other, thereby enabling cooperation in preventing and detecting fraud and money laundering. Therefore a measure of the privacy risk associated with the deployed privacy-preserving technique is becoming increasingly important. For use within individual organizations, in addition to privacy, it would be desirable to preserve forensics, i.e. that actions performed by the detect-and-react subsystems in response to fraud detections should also enable traceability and information for revealing the nature of the discovered frauds and responsible actors.

10.2. Testing, Certification, and Standards Alignment

To instill confidence in the safety and security of agentic fraud detection systems, rigorous testing procedures, independent certifications, and compliance with widely accepted standards are essential. Safety-critical systems are typically subjected to extensive validation against relevant specifications but lack established industry standards. Therefore, dedicated evaluation protocols are necessary. Third-party certifications, such as compliance with regulations governing operation in sensitive domains, also enhance trust. Testing methodologies, conformance to regulatory frameworks, and certification processes for cloud environments reinforce the safety, reliability, and risk management aspects of agentic systems.

Validation strategies must encompass evaluation of monitoring components for potential false positives, trade-off analyses, auditing and testing of the detection mechanisms, and deterrence studies. Extensive reviews of patterns for AI-based systems provide valuable references for such protocols. The Trustworthy AI Assessment List and auditing frameworks for ethical and responsible AI, AI risk management, and privacy protection further inform validation requirements. Security-focused testing recommendations and hardware secure element validation guidelines also inform the necessary conformance checks for fraud detection systems.

11. Case Studies and Practical Implications

Deployment scenarios illustrate real-world applicability and yield insights for others embarking upon similar initiatives. Distinct requirements and design choices characterize agentic-fraud-detection systems in personal and commercial lines. Detection increasingly relies on open-source frameworks that emphasize flexibility, scalability, and resiliency at every layer of the architecture. Governance considerations center on oversight of data-subject consent and control, safe cross-border transfers, and careful scrutiny of privacy-preserving techniques.



Organizational and stakeholder risks must be assessed and managed throughout the process, from initial project exploration to operationalization of the agentic decision system.

Cloud-native insurance platforms are especially well suited to the agentic paradigm in personal lines. The rapidly changing nature of agentic-pandemic-related fraud lends itself to the flexibility and data accumulation typical of cloud-native ecosystems. Ground-up development and deployment in such environments, driven by considerations of scale, resilience, and practical limitations of monolithic architectures, reduce the risk of design flaws. In contrast, agentic systems in commercial lines often build on established investment in other decision domains, as is the case for a fraud detection capability meeting the specialized subtleties introduced by this risk area.

Table 4: Cloud-Native Architecture Components for Fraud Detection

Component Name	Function	Importance
Data Ingestion Layer	Collects data from sources	High
Data Lineage System	Tracks data origin and flow	High
Feature Engineering	Prepares input features	High
Agentic Scoring Engine	Calculates fraud probability	Critical
Monitoring System	Tracks performance and drift	Critical
Decision Engine	Flags fraud/non-fraud	Critical
Audit & Logging	Stores decision history	High
Feedback Loop	Improves system over time	Critical

11.1. Deployment Scenarios in Personal Lines and Commercial Lines

Fraud detection in personal lines differs from that in commercial lines of the insurance sector. Insurance products such as motor or household insurance are of lower financial value and involve low-frequency transactions. Anomalies are rare, and insured members usually have little prior involvement with the company. Consequently, data used for anomaly detection is often insufficient to build a reliable model. On the other hand, in commercial lines, data volume increases and specific methods such as those based on agentic AI can be applied.



A cloud-native insurance platform with agentic capabilities for real-time fraud detection in personal insurance lines is a proven deployment for processing voluminous data streams with a structured and reliable infrastructure. Differences in privacy considerations surrounding personal and commercial lines have a significant impact on the deployment of an agentic fraud detection system for commercial lines. The potential effect on the fraud detection model's decision-making process is examined, including privacy-preserving techniques such as data anonymization, enabled by publicly available databases. These measures bolster stakeholders' confidence in adopting an agentic AI engine for fraud detection, allowing them to focus on business expansion rather than on security concerns.

Equation 5. Recall / True Positive Rate

Actual fraud cases:

$$TP + FN$$

Detected fraud cases:

$$TP$$

Therefore:

$$Recall = TPR = \frac{TP}{TP + FN}$$


11.2. Risk Assessment and Change Management

Commencing agentic AI adoption with fraud detection in personal lines reduces complexity and exposure; other change-management aspects remain pertinent. ID Care in Taiwan employed eKYC to streamline onboarding, improving monitoring and shortening fraud case-discovery times. Security for Business mapped risks from unauthorized access, data breaches, compliance failure, and vendor service stoppage; engaged users and customers; and established detection and response processes.

Agentic AI enables an enterprise-wide agility response to changing conditions and demands. Enterprises need to mitigate exposures while enhancing their ability to change and grow; agentic systems provide the patrol functions that search for adaptive change opportunities



and assess risk/return trade-offs. Careful evaluation of agentic-technology decisions helps calibrate future likelihood of change failure. Stakeholders need to understand potential future outcomes and agentic-technology investment consequences. Kwan Chai set up a personal-assets security system to manage assets, records, registration, and contact notifications.



12. Results

Core findings are summarized within a consolidated framework, accompanied by a high-level overview of relevant data ecosystems and associated interoperability requirements.

Data ecosystems constitute a critical dimension of agentic fraud detection, supporting both detection continuity and technology validation while facilitating broader analytical and operational capabilities. Building and maintaining a data ecosystem necessitates close attention to matters of governance, especially related to data provenance, privacy, compliance, and quality. The considerations identified in this study inform the design of such an ecosystem and, more broadly, help steer AI adoption in risk-aware directions.

Development and deployment of agentic systems entail a variety of risk considerations beyond those related to data quality. Trade-offs between detection effectiveness and user experience, the need for external validation, emerging capabilities for reverse engineering and



bias detection, and the process of deploying agentic systems in production are all relevant to decisions shaping the wider adoption of agentic fraud detection.

Size & Share 2025-2032 (USD Billion)




Fig 6: Agentic AI In Prevention & Fraud Detection Market

13. Conclusion

Framed as a policy-ready synthesis for insurance practitioners and stakeholders overlaid with scholarly methodologies, the research examines the implementation of Agentic AI (with self-learning, adaptive, and real-time responsiveness capabilities) in cloud-native insurance Fraud Detection Systems (FDS), the design of Fraud Detection Models (FDM) that dynamically replace Rule Sets (RS) within these FDS, and the methodologies for continuously training, monitoring, and validating these FDM. Such a design pattern empowers fraud detection and prevention in insurance values chains. The practical exploration model provides a flexible-enough framework for insurance experts tasked with FDS management, who may have limited knowledge of Machine Learning yet need to explore and validate the deployment of Agentic Fraud Detection FDMs. It can be adapted and fine-tuned according to the specific dimensions of any use case or context of interest, generating practical insights related to the deployment and use of Agentic AI to these insurance FDS.



Although the formal framing can be better suited and complemented with diverse discussions in Machine Learning, Agentic AI, and Normative AI, the research can be very revealing in academic explorations focusing on FDS. The proposed Insurance-FDS deployment model summarizes the core applicability domains: it outlines the relevant Data Ecosystems; navigates the Privacy, Security, and Compliance issues; discusses the different methodologies for Fraud Detection These Detect Reasoning Method homes in on Model Architectures for Real-Time Fraud Scoring—embodying the technological aspects for FDM deployment—before addressing Trust, Ethics, and Governance of Autonomous Insurance Systems and proposing Deployment Scenarios relevant for the Design and Validation of Insurance FDS Expert And End-User Consulted.

References

1. Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. *PeerJ Computer Science*, 9, e1278.
2. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.
3. Bhandari, P., Prithviraj, J., Saxena, R., Kumar, D., Mangala, N., & Chourasiya, P. (2026). Harris Hawks Optimization–Based Deep Learning for Resource-Aware Cloud–Fog Computing. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0. <https://doi.org/10.1109/otcon69332.2026.11629982>
4. Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University—Computer and Information Sciences*, 35(1), 145–174.
5. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: A review. *Frontiers of Computer Science*, 19(9), Article 199609.
6. Avinash, K., Ranga Reddy, V. A., Manohari, D., Kothuri, S., Geetha, R., & Swarna, B. (2026). Multi-Source Feature Integration for Short-Term Load Forecasting in Off-Grid Residential Microgrids Using Stacked LSTM Networks. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 501–505). IEEE. 2026 14th International



- Conference on Smart Grid (icSmartGrid).
<https://doi.org/10.1109/icsmartgrid70882.2026.11680767>
7. Debener, J., Heinke, V., & Kriebel, J. (2023). Detecting insurance fraud using supervised and unsupervised machine learning. *Journal of Risk and Insurance*, 90(3), 743–768.
 8. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, Article 116429.
 9. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, Article 122156.
 10. Seenu, A., Kolla, S. H., Aitha, A. R., Amistapuram, K., Abhireddy, N., & Yandamuri, U. S. (2026, June). Agentic Enterprise Intelligence: A Unified Architecture for Secure RAG-Enabled LLM Automation in Insurance Risk and Fraud Ecosystems. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-5). IEEE.
 11. Sulaiman, R. B., Schetinin, V., & Sant, P. (2022). Review of machine learning approach on credit card fraud detection. *Human-Centric Intelligent Systems*, 2, 55–68.
 12. Pourhabibi, T., Ong, K.-L., Kam, B. H., & Boo, Y. L. (2020). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*, 133, Article 113303.
 13. Rani, S., Thavara, S. S., Kr, A., Naguband, A. R., Garg, A., & Vajpayee, A. Financialization of Sustainability: ESG Signalling, Market Valuation and the New Economics of Corporate Legitimacy.
 14. Autonomous credit card fraud detection using machine learning approach. (2022). *Computers & Electrical Engineering*, 102, Article 108132.
 15. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
 16. Singh, K. U., Gupta, S. K., Dubey, Y., & Mangalampalli, B. M. (2026, June). Enhancing Healthcare Cyber Defense with Machine Learning-Based Attack Prediction and Prevention. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
 17. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask,



- A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, Article 119.
18. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60.
 19. Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry. (2026). *Journal of European Economic History*, 7(2), 68–88.
<https://doi.org/10.61336/jeeh/26-2-7> (Original work published as Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry)
 20. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, Article 122156.
 21. Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. *PeerJ Computer Science*, 9, e1278.
 22. Anand, V., A, S. K., Vemulapalli, N., Loganathan, R., Guntupalli, A., & Ahammed, S. R. (2026). Federated Deep Learning for Decentralized Cybersecurity in Smart Grid Communication Networks. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 475–481). IEEE. 2026 14th International Conference on Smart Grid (icSmartGrid). <https://doi.org/10.1109/icsmartgrid70882.2026.11680752>
 23. Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University—Computer and Information Sciences*, 35(1), 145–174.
 24. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, Article 116429.
 25. Babu, A. J., Yandamuri, U. S., Recharla, M., Pamisetty, A., Goli, M., & Kolla, S. H. (2026, June). Data Analytics and AI Solutions for Digital Transformation in Hospitality and AgriTech Industries. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-7). IEEE.
 26. Sulaiman, R. B., Schetinin, V., & Sant, P. (2022). Review of machine learning approach on credit card fraud detection. *Human-Centric Intelligent Systems*, 2, 55–68.
 27. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.



28. Mahadevan, S. (2026). Governed Serverless Automation Ecosystem for AI-Driven Enterprise Integration and Sustainable Cloud Operations. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(16s), 1561-1570.
29. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
30. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3, Article 119.
31. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on “Predicting self-image satisfaction after adult spinal deformity surgery: a machine learning approach using patient phenotypes.” *Spine Deformity*. <https://doi.org/10.1007/s43390-026-01400-3>
32. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60.
33. Debener, J., Heinke, V., & Kriebel, J. (2023). Detecting insurance fraud using supervised and unsupervised machine learning. *Journal of Risk and Insurance*, 90(3), 743–768.
34. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: A review. *Frontiers of Computer Science*, 19(9), Article 199609.
35. Raut, B., Garg, A., Mangalampalli, B. M., Jindal, N., Mahajan, V., & Saraswati, K. (2026). Multidisciplinary Strategies for Improving Maternal and Reproductive Health Outcomes and Mental Outcomes. *International Journal of Drug Delivery Technology*, 16(34s). <https://doi.org/10.25258/ijddt.16.34s.36>
36. Pourhabibi, T., Ong, K.-L., Kam, B. H., & Boo, Y. L. (2020). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*, 133, Article 113303.
37. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, Article 122156.
38. Sulaiman, R. B., Schetinin, V., & Sant, P. (2022). Review of machine learning approach on credit card fraud detection. *Human-Centric Intelligent Systems*, 2, 55–68.
39. Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. *PeerJ Computer Science*, 9, e1278.



40. Nemani, R., Singuluri, I., Ramani, V. S., Guntupalli, A., Mopathi, U., & Madineni, A. V. S. (2026, July). A Machine Learning-Augmented Penalty Difference Approach for Transportation Cost Optimization. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 511-516). IEEE.
41. Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University—Computer and Information Sciences*, 35(1), 145–174.
42. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, Article 116429.