




Self-Adapting AI Agents for Financial Crime Surveillance

1st Srinivasa rao challa , 2nd David Schneider

Srinivasa Rao Challa, Sr. Manager, Charles Schwab , srinivas.r.challa.sm@gmail.com

David Schneider , Research Assistant

Abstract

Financial crime and the efforts to combat it sit at the intersection of fear and greed, the two most powerful drivers of human behavior. The gap between these motivational extremes shapes the risk borne by every party in an economic exchange and can ultimately undermine the integrity of the financial system itself. Regulatory authorities oversee how supervised entities work to prevent and counter financial crime, with the goal of safeguarding the stability of domestic and international financial systems. Yet compliance is frequently viewed—by regulated entities and other market participants alike—as, at best, a necessary burden and, at worst, an obstacle to business. As a result, supervised entities continue to struggle with making compliance both effective and efficient.

A collaborative multi-agent approach offers a way to reframe financial crime compliance: from a regulatory obligation alone to a genuine driver of business value for organizations and their broader ecosystems. Such systems can dynamically direct resources toward the alerts that matter most, sharpening situational awareness for governance and control. By embedding an understanding of risk labels—including evasive behavior patterns—directly into detection logic in a distributed way, these systems achieve compliance by design, without relying on data scientists to hand-craft rules.

keywords: Adaptive Multi-Agent Artificial Intelligence Systems, Financial Crime Compliance, Risk Prediction-Crime Detection, Financial Crime, Artificial Intelligence, Machine Learning, Explainable Artificial Intelligence.

1. Introduction

The financial ecosystem can be seen as a complex system where entities belonging to different business domains interact, exchanging resources, and providing services to each other. These complex interactions create a wide range of risks of inappropriate behaviors among financial institutions and participants. Financial crime affects the industry as a whole and, therefore, the supervisory authorities of each country create regulations that require financial institutions, as well as other entities within the financial ecosystem, to monitor their transactions and the status of their clients in search of such inappropriate behaviors. This action is called compliance, defined as the process of ensuring that the set of rules established for a financial environment are being followed.

A financial institution or service provider, regardless of its business model, act as the first line of defense against inappropriate conduct and must, therefore, have a Compliance structure that monitors the transactions and situations of its customers. Monitoring operations, however, is a very expensive task and compliance areas need to have people, rules, and technologies that

support them in their mission. Technologies with a focus on Automation and Artificial Intelligence (AI) have emerged to support compliance with these regulations. For instance, machine learning can automate the detection of events that may indicate money laundering (ML). Machine learning models can also monitor social networks and news about people and organizations with relationships in financial ecosystems, generating warnings of latent risks.

1.1. Background and Significance

Financial crime manifests itself in many forms and places a significant burden on society. For example, the World Economic Forum estimated in 2020 that financial crime cost the global economy USD 1.4 Trillion, equivalent to roughly 1.5% of global economic output. The extraordinary costs associated with prevention, investigation, prosecution and branding damages, along with the huge efforts dedicated for anti-money laundering by law-enforcement and HMRC, directly contributed to the recent call for a rethink of the current paradigm. Major regulatory breaches can lead to fines running into billions of dollars and erode confidence in the financial system.

This is particularly relevant given the financial institution’s efforts to act as a gatekeeper to the financial system, to prevent and detect financial crime and support law enforcement and regulators. Financial institutions are required to spend millions on compliance and yet the success of the current paradigm is far from proven. The goal of Adaptive Financial Crime Compliance is to advance the current paradigm, enabling compliance regimes to be responsive, efficient, effective, and focussed on areas of real risk.





Fig 1: Multi-Agent AI Systems for Adaptive Financial Crime



1.2. Research design

Research design addresses regulatory compliance and risk prediction through Multi-Agent System (MAS) technologies capable of adaptively satisfying formal compliance requirements while also predicting financial crimes at run-time. Three research directions are identified. First, the MAS framework supports real-time scoring of actors in the financial ecosystem by predicting a range of risk indicators from transaction data using supervised learning models. Second, stability and robustness of these real-time models are considered, with a focus on the quality and availability of the underlying risk assessment data. Third, the MAS framework addresses the adaptive enforcement of compliance rules, especially when operating in unfamiliar contexts or under extraordinary conditions, while satisfying explainability and auditability requirements.

The research also identifies open issues. Responding to volatility in crime patterns or unstable run-time models requires either continuous real-time retraining on fresh data or a transfer learning strategy. As operational latency can increase with the number of rule-specific agents, it may not be feasible to support a large number of expert agents. Further, although the detect-and-react paradigm can be effective and scalable, it raises trust issues if detection is unreliable, unsynchronized, or not supported by credible laboratory testing. The distributed nature of knowledge and policy management helps with compliance by design, yet the quality and provenance of information remain critical concerns.

Equation 1: Real-time risk score equation

Step 1: Start from contribution of each feature

Each feature contributes:

$$w_1x_1, w_2x_2, w_3x_3, w_4x_4$$

Step 2: Add all feature contributions

Total raw risk:

$$R_{\text{raw}} = w_1x_1 + w_2x_2 + w_3x_3 + w_4x_4$$


Step 3: Add intercept/bias term

A baseline system risk is usually added:

$$R = b + w_1x_1 + w_2x_2 + w_3x_3 + w_4x_4$$

Final equation

$$R = b + \sum_{i=1}^4 w_i x_i$$



2. The Landscape of Financial Crime Compliance

Financial crime is a complex social phenomenon involving a wide range of acts. Financial institutions and businesses are obliged to comply with anti-money laundering (AML), counter-terrorism financing (CTF), and other anti-financial crime (AFC) regulations to prevent such crime and its consequences. Recent developments indicate a growing adaptation of regulations and the adoption of a risk-based approach in AFC compliance. However, regulatory approaches are behind current trends and exploitation of distributed systems and emerging technologies such as blockchain and cryptocurrencies. In addition, gaps and blind spots in compliance systems still exist. Recent developments at regulatory authorities and international agencies indicate trust in financial and fintech ecosystems that emulate and absorb necessary functions to offer trustless transactions.


Financial institutions and businesses face a complex web of legal obligations in how they conduct operations. Compliance costs continue to decimate revenues, with regulatory authorities unable to demonstrate a reduction in financial crime through application of the regulations. Nevertheless, the financial services sector is immersed in a cat-and-mouse game with criminal forces; hence spending is focused on the task of ensuring freedom from legal actions for bank boards, businesses, and government ministers. In the particular case of AML, the cat is winning the game, with nearly \$1 trillion laundered every year.

Component	Purpose	Main Functions	Expected Benefit
Real-time risk prediction agents	Predict financial crime risk across actors and transactions	Score entities using transaction data, historical patterns, and external indicators	Early identification of suspicious behavior and dynamic risk prioritization
Compliance monitoring agents	Enforce regulatory and operational compliance requirements	Monitor rule adherence, detect compliance breaches, trigger alerts	Stronger compliance coverage and reduced manual oversight burden
Situational awareness agents	Detect abnormal ecosystem behavior and provide context	Observe cross-entity activity, identify anomalies, aggregate signals from other agents	Better contextual understanding of suspicious events
Rule management and policy agents	Maintain adaptive compliance rules in distributed form	Translate expert-defined rules into machine-processable representations, update policies, coordinate enforcement	Greater adaptability to evolving crime patterns and regulations

Table 1: Core Components of the Proposed Multi-Agent AI Framework

2.1. Regulatory Frameworks and Compliance Objectives

Financial crime compliance is a multifaceted concern encompassing diverse securities and regulatory frameworks across banking and insurance supervision, stock exchanges, money laundering and terrorist financing prevention, data protection, anti-fraud, anti-corruption, and sanctions. The variety of separate compliance regimes increases complexity and costs. For instance, Anti-



Money Laundering (AML) compliance is mandated in most jurisdictions under set rules, such as the Financial Action Task Force special recommendations. These rules create massive costs for businesses. The European Union has also begun a shift toward a single compliance regime incorporating the aforementioned areas. The European Legislative Proposal for Money Laundering Prevention proposes the establishment of a single supervisory authority for AML supervision, capable of establishing powers and procedures at union level for the prevention of money laundering. However, for the moment, such compliance is highly fragmented, both from a jurisdictional and from a Business Aspect view.


Moreover, compliance efforts are often seen as a necessary cost—but not as an opportunity to enhance business resilience and success. Traditional approaches typically require separate teams in each jurisdiction and area, generating redundant costs and potential inefficiencies stemming from limited situational awareness across the whole enterprise. A multi-agent AI ecosystem could serve as the foundation for an adaptive, flexible, and intelligent compliance engine that rethinks and repositions compliance as a twofold opportunity: to provide better and more secure services and to control costs by automating, simplifying, and streamlining compliance with total compliance-proof conditions.



Fig 2: Regulatory Compliance

2.2. Limitations of Traditional Compliance Approaches

Despite substantial investments in financial crime compliance, criminals continue to exploit the financial ecosystem. There are three principal reasons for this continuing failure. First, compliance departments are stymied in their objective to detect financial crime because of the overwhelming volume of low-quality alerts generated by traditional detection systems. This alert burden is intense as an estimated 95% of alerts are false positives. Second, compliance departments are frequently unable to proactively detect the latest crime trends: it is estimated that only 5% of the alerts raised concern high-risk clients, such as the Breitling group that laundered illicit funds for global drug cartels. Finally, suspicious activity report filings have plummeted in recent years amid the rising numbers of alerts; in 2019, the number of SARs for 1000 transactions was approximately one-fifth the level in 2002.



Traditional financial crime compliance relies on a pragmatic combination of time, money and resources. The general approach is to meet compliance obligations, but avoid failures that attract the attention of regulatory authorities. These three factors combine to create a new opportunity for criminals to perpetrate heinous acts from the comfort of their armchairs in countries that do not respect the rule of law or engage in proactive investigations into financial crime. The multi-agent system framework addresses this issue by supporting demonstrably effective compliance that can be audited by regulators for compliance governance.

Equation 2: Status-conditional probability equation

Let:

- S_{t+1} = future status at next time step
- R_t = current risk score
- X_t = other current observed features/events

We want:

$$P(S_{t+1} | R_t, X_t)$$

Step 1: Start with conditional probability definition

Using Bayes' theorem:

$$P(S_{t+1} | R_t, X_t) = \frac{P(R_t, X_t | S_{t+1})P(S_{t+1})}{P(R_t, X_t)}$$


Step 2: Interpret each term

- $P(S_{t+1})$: prior probability of status change
- $P(R_t, X_t | S_{t+1})$: likelihood of observing current risk/features if status change happens
- $P(R_t, X_t)$: normalization term

Step 3: Write the final status-conditional model

$$P(S_{t+1} | R_t, X_t) = \frac{P(R_t, X_t | S_{t+1})P(S_{t+1})}{P(R_t, X_t)}$$

3. Multi-Agent Systems: Concepts and Rationale



The term agent refers to an abstract entity that is capable of perceiving and acting upon its environment. Multi-agent systems (MAS) consist of multiple interacting agents, capable of communicating and jointly working towards a common goal or set of goals. The multi-agent paradigm supports decentralization, which helps overcome scalability issues, as well as flexibility and resilience in dynamic environments. MAS are often classified into two groups: homogeneous systems, where agents share the same architecture; and heterogeneous systems, where a different architecture is adopted for each agent type. MAS may also be classified according to the coordination mechanisms employed, including reactive, deliberative, hybrid, contract net, and market-based systems. In reactive systems, interaction is typically governed by rules that define the agents' behavior. Deliberative systems rely on a more complex model of the world, based on beliefs, desires, and intentions. Hybrid systems combine elements of both paradigms. In contract-net systems, agents announce their capabilities and request services, while market-based systems rely on bidding behaviors.

Regulatory frameworks impose strict obligations, especially on large financial institutions. Compliance of individual entities is therefore essential for a healthy financial ecosystem. Such obligations include the prevention and detection of money laundering and terrorist financing, the combatting of fraud and bribery, and the reporting of suspicious activities. However, as financial services have migrated online and into the cloud, and become heavily automated and interconnected, traditional compliance systems have become more complex, centralized, and expensive. Process-oriented information technology platforms, typically deployed only by large financial institutions, offer only limited support for cost-effective compliance across payment services, securities services, banking sectors, and examiners. A multi-agent system architecture allows independent financial services providers to share information and create a situational awareness of compliance-relevant events across the entire financial ecosystem.

Modern regulatory frameworks place significant compliance burdens on financial institutions, requiring robust mechanisms to prevent money laundering, detect fraud, combat bribery, and report suspicious activities. As financial services increasingly operate in digital, cloud-based, and highly interconnected environments, traditional compliance systems have struggled to keep pace due to their centralized, complex, and costly nature. These legacy, process-oriented platforms—often accessible only to large institutions—offer limited flexibility and scalability for smaller or independent entities. In contrast, a multi-agent system architecture presents a more adaptive and collaborative approach, enabling diverse financial service providers to operate as interconnected yet autonomous agents. By securely sharing relevant data and insights, such systems can enhance real-time situational awareness, improve the detection of compliance risks, and foster a more resilient and transparent financial ecosystem overall.





Fig 3: Multi-Agent Systems

3.1. Agent Architectures and Coordination Mechanisms

Multi-Agent AI systems for adaptive financial crime compliance and real-time risk prediction; maintain an objective, evidence-based, formally structured academic tone throughout.

Agent-based solutions traditionally exploit mechanisms such as communication and negotiation to instantiate cooperative, altruistic or self-interested socially aware agents working for a common goal. Agent types range from the most simplistic — often representing a single entity through events and cues — to the most expressive — where reasoning is coupled with belief modelling — allowing for collective agent behaviour. In recent explorations, collective agent functions have addressed the act of governing agents; regulating agent behaviour and coordinating across multiple agents are essential requirements of any multi-agent platform designed to address a complex problem. Governance mechanisms regulate communication, contract or crime type, risk appetite, and evaluation criteria; they set the rules of the game and define the agent ecosystem.

In high-stakes, real-time systems such as financial crime compliance, speed is paramount. Responses to dynamic events should therefore be automatic and fast; adaptivity to changes in policies and risk appetite should require little manual intervention; and cooperation between agents should require minimal — if any — supervisory overhead. Desired levels of system adaptivity range from a basic capability to detect and trigger the execution of an alternative strategy — such as re-routing a transaction to another entity — to the auto-generation of new strategies through the analysis of recent detected events, their consequences, and the roles played by the involved agents.



3.2. Adaptivity and Real-Time Collaboration

Difficulties associated with applying AI solutions to compliance and monitoring tasks have led to a preference for traditional models and systems, despite the appeal of proactive solutions. AI can contribute in two ways. Multiple, purpose-specific AI agents dynamically cooperate to monitor regulatory compliance in financial institutions, including the adaptation of rules to shifts or anomalies in transaction or behavior patterns. Real-time risk evaluation in complete financial systems is a second use case. This distinct capability exploits historical data from problems, incidents, or major changes in the ecosystem (like the Russian invasion of Ukraine or the COVID-19 pandemic) to evaluate operators' riskiness on a permanent basis. Both strategies focus on financial crime.

Detecting cybercrime and mitigating its impact is challenging due to the often-lax security of parties involved in electronic payment transactions. In this sphere, the built environment could have an important role in managing and actively tackling financial cyberattacks. In recent years, several prominent market players have been accused of providing financial services to sanctioned countries or persons. Monitoring and sanctioning corresponding breaches is a difficult exercise as the category of risk can only be assessed ex-post. State-of-the-art machine-learning techniques applied to a 4-year time series of deposits and WASH transactions of the 1,700 financial institutions partially covered by the Financial Action Task Force could produce an accurate risk label for those institutions. Integrating this model in a multi-agent AI system could therefore prevent the occurrence of such serious incidents.

Equation 3: Anomaly detection rule

Let:

- x = current transaction/activity vector
- μ = mean normal behavior
- σ = standard deviation of normal behavior
- $A(x)$ = anomaly score
- τ = alert threshold

Step 1: Measure deviation from normal behavior


For one variable, standardized deviation is:

$$A(x) = \frac{|x - \mu|}{\sigma}$$

Step 2: Compare deviation with threshold

If deviation is large enough, trigger anomaly:

$$A(x) > \tau$$



Step 3: Write decision rule

$$\text{Alert} = \begin{cases} 1, & A(x) > \tau \\ 0, & A(x) \leq \tau \end{cases}$$

4. Real-Time Risk Prediction in Financial Ecosystems

Timely forecasting of aberrant behavior manifestations within financial ecosystems is vital for minimizing operational risk and safeguarding both business and customer interests. Risk prediction is, hence, a crucial prerequisite for implementing compliance mechanisms, facilitating the real-time detection of behavioral deviations by normal agents. The majority of established approaches rely on off-line predictive modeling processes that typically require long time horizons and abundant domain expertise for the manual definition of safety, security, and compliance concerns. This severely restricts their use in cases where changes in the underlying systems are frequent and can entail significant shifts in the nature of anticipated risks.

Timely prediction of risk situations necessitates an alternative modeling technique that produces up-to-date risk scores for the analytical functions of the operational monitoring agents in an adaptive multi-agent compliance system. Financial and non-financial incidents associated with identified or anticipated risk situations must be monitored through external indicators, such as data and reports from standardization, supervision, and regulatory entities, declared sanctions lists, police crime reports, forecasted macroeconomic variables, and any other public or web-sourced external data relevant to the implementation of specific banking rulebooks. These sources can then be exploited for risk-feature engineering. Assessing specific, novel data sources pointed out in the literature may further increase the efficiency of risk prediction.





Fig 4: Real-Time Risk Prediction in Financial Ecosystems

4.1. Data Sources and Feature Engineering

The nearly three billion daily transactions in the UK alone require an immense human and financial effort to detect fraudulent activities. In general, enterprises, especially banks, apply anti-money laundering, counterterrorist financing, and combating the financing of proliferation policies and programs to satisfy various national and supranational regulations. Russia, for example, has implemented specialized requirements and applied hardware-computing resources for three large banks: VTB, Sberbank, and Gazprombank. Nevertheless, the costs of compliance remain high and collapse proportions with little or no outcome in flagging suspicious activity, transactions, or endeavors, and no consequences after these actions of tendency.

Commonly, three main categories of data sources exist for all enterprises: internal open, proprietary data; internal—and moderately private, if necessary—sensitive data; and external sources. Like any traditional machine learning task, defining the possible feature space or projection of each data object is essential. Each enterprise might have major roles and transactions supporting uncertain systemic risk. Classification and analysis of the results from central banks and supervisory agencies offer possibly significant features. Systematic, unsupervised clustering applies to each detected class to seek microcategories or substructures with a high average cumulative distribution for the detected task, and topological properties are checked according to the desired microtask. The feature space is reduced to include only the transactions that persistently are part of the same base group. Advanced agents may have label-helping tasks that are supervised or assisted.



4.2. Modeling Techniques for Timely Risk Scoring

Four modeling families are considered for operational risk prediction: status-agnostic, status-conditional, supervised, and unsupervised. Status-agnostic models receive historical time-series data over a radius of one day and output a risk score without knowing the status of the target entity. Status-conditional models, on the other hand, receive a risk score per entity as input and predict the probability of a status update within the next hour. Supervised models leverage historical changes in the status variable as labels. Unsupervised models generate risk scores and labels purely by estimating submodes over the time series.


Timely operational risk scores are typically aggregated at a higher time resolution than that of the underlying status change events. As such, status-conditional models are often applied to predict changes in status category within the next hour. However, as recent events may impact the likelihood of a status change and current risk scores provide valuable information, a Bayesian-like model is expected to yield the best performance: a status-conditional model that incorporates the risk score into its feature set. As the data naturally lends itself to the supervised learning paradigm and labels are generated dynamically with every incoming event, supervised models are also considered for such predictions.

Dimension	Description from Article	Operational Role	Challenges Identified
Internal open/proprietary data	Enterprise-owned operational and transaction data	Baseline source for transaction monitoring and institutional behavior analysis	Scale, heterogeneity, and integration difficulty
Internal sensitive data	Moderately private institutional/customer data	Supports more precise compliance and anomaly detection tasks	Privacy, access control, and regulatory constraints
External data sources	Supervisory reports, sanctions lists, police reports, macroeconomic indicators, public/web data	Enriches risk-feature engineering and ecosystem-wide situational awareness	Data quality, provenance, timeliness
Status-agnostic models	Predict risk score without current status label as input	Useful for broad entity-level risk estimation from historical time series	Lower contextual specificity

Table 2: Risk Prediction and Data Modeling Structure

4.3. Evaluation Metrics and Operational Considerations

Regardless of the modeling techniques employed, practical solution development demands careful consideration of evaluation metrics that properly reflect model performance in terms of real-world banking, payment, or security transactions. Different financial crime applications entail detecting diverse risk classes for banking institutions: A money laundering process is typically initiated by a low-risk profile crime and concludes with a high-risk profile crime. The initial crime is generally indicated by a suspicious transaction report (STR) filed by another institution, although STRs against non-compliant crimes represent only a



small percentage of the total crime volume. By contrast, payments support retail crime directly, being an integral part of each. Thus, judges and policy makers must be the direct victims of an abuse of an innocent money-sending application or service. For payments, early warning systems should ideally try to classify anomalous-use-of-service alerts created by the local exchange risk engine. These are categorized as hi-risk alert-do-not-serve-risk profiles defined by a host-service usage-facial mortgage policy. Thus, an implicit early warning-enumerated target-stack for each transaction should also be created.

While penalizing false negatives in these cases of direct victimization makes sense for some problems, the problem of credit cards is more of a public-policing nature. Anomalous credit-card usage generally indicates a private signal for public crime and transit-facility anomalous use of an anonymous infrastructure. Here, public sector sponsors should care more about detecting plausible terrorist-planning chains. In this situation, risks can be tuned on balance, steered away from annoyance to exploration. A further point is that to keep honesty or privacy in transit, attacks or disturbances must also be detected. These services make no direct profit but permit confidence in anonymity by criminal groups. Thus, these classes of SCUs need more than the above services: A complementary credit alerts system should be in place, scoring the plausibility of association with a large number of operational-users populations. This further credit-alert class is discreet since it merely keeps track of a different kind of population, making sure that planning groups attack or commit other non-fatal acts of crime.

5. Adaptive Compliance through Multi-Agent Collaboration

The supervision of financial crime compliance typically relies on rule-based expert systems that distribute rules across cooperating world models. These rule bases, however, are not directly managed by the systems themselves but rather externally engineered and maintained in a gated static manner, much like manually crafted compliance workflows. Such an approach does not offer the adaptability, up-to-date situational awareness, or fitted explainability required and recommended in current regulation. A multi-agent architecture can fulfill these roles, offering analytical capabilities for situational awareness to detect inconsistency and flagging of sensing data that might need closer investigation. Anomaly detection can be used to identify and issue alerts for supported and new types of issues alike.

Compliance requires situational awareness and agility to adapt to new threats and evaders, enabling timely mitigation of the next breach or episode. This can be achieved through a suitable multi-agent architecture that enables internally-provided adaptivity and agility. The architecture operates over three main dimensions. Instrumental functions (process monitoring, rules) reside externally in the knowledge base and are by design governed and maintained at a higher security clear level. They undertake operational roles needed for compliance (monitoring, control) and protection (prevention, detection). No externally engineered rule is executed without first being flagged by a respective internal entity. The latter issue alerts for changes in supported or newly detected categories of issues, collectively termed "anomalies."




Fig 5: Overview of our multi-agent collaboration platform

5.1. Distributed Rule Management and Policy Enforcement


The modular architecture of a multi-agent system (MAS) promotes the creation of specialized agents with narrowly defined tasks, enabling the development of distributed subsystems for rule management, detection, and mitigation of financial crimes such as money laundering, financing of terrorism, and bribery. A particular focus is given to the clear representation and interpretation of operation rules for human users. These rules, initially authored by expert users in natural language, are transformed into a machine-processable form, and supplemented by a well-defined ontology. Rule-based systems for anomaly detection are continuously updated and the underlying data is actively monitored by meta-layers ensuring that newly discovered situations and newly acquired ability to handle them are successfully integrated into all relevant detection systems. Achieved detections are translated into appropriate alerts according to situational parameters, generating Information that is presented, for validation and approval, to qualified incident response teams. Possible remedial actions are standardised according to the nature of the incident and their description is also made understandable to users that are not experts in the field.

Regulatory Compliance is ensured through the support of specialised agents, controlling and documenting the operation of systems. Self-disclosure and compliance-by-design recommendations are generated by distilling operational data into understandable form. Compliance expertise is formalised in operational Design Codes for adaptive use. Audit Trails and Explanation Agents assist accountability through situationally appropriate operation narratives in end-to-end natural language.

Equation 4: Precision equation

Let:

- TP = true positives
- FP = false positives



Step 1: Define all predicted positives

All positive alerts generated by the system are:

$$TP + FP$$

Step 2: Find the fraction that are correct

Precision is the fraction of predicted positives that are truly positive:

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

Final equation

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


5.2. Anomaly Detection and Situational Awareness

To provide real-time situational awareness for financial crime compliance, a dedicated agent monitors the entire ecosystem and identifies anomalies in the activities of participants. Situational awareness is commonly understood as "the perception of environmental elements within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future". Like other participants in the ecosystem, the situational awareness agent is able to assess the normality status of its environment and detect any unusual behaviour using anomaly detection techniques. For instance, it can identify whether the activity of a particular bank, payment processor, or crypto asset exchange is consistent with its normal activity pattern. In that respect, the agent simulates an expert in the fight against money laundering and terrorist financing, producing information that is useful for the participants' compliance managers as well as for law enforcement agencies.

In addition to identifying anomalous behaviour, the situational awareness agent must provide context by explaining why a given event is determined to be anomalous. One way to achieve this is by aggregating information from the other agents about the transaction party, the transaction structure, or the transaction purpose, for example about all participants involved in a transaction at the same time, as detected by the adjacency agents. Proper parametrisation of the anomaly detection techniques is crucial, especially in such an application context where the distribution of the underlying processes may change suddenly and dramatically. The situational awareness agent should also offer explanations, helping the users to understand the reasons of an anomaly. This is particularly important for compliance staff and law enforcement agencies, which require substantiated rationale when blocking or investigating transactions.

5.3. Explainability, Auditing, and Governance

A major hurdle for the practical implementation of financial crime detection and prevention systems is their acceptance and trust from law enforcement, judicial bodies and the involved parties. These systems usually recommend actions for the responsible



entities, rather than executing them autonomously. These recommendations are often not followed due to the lack of explainability. A multi-agent approach, where agent actions and decisions are inherently interpretable insights for the encompassed tasks, might circumvent this issue. The self-organizing nature of the collaboration enables the agents to focus on the distribution of the financial crime detection processes – rather than the distribution of roles in every single situation – boosting their capacity to jointly respond to significant threats.

Another major requirement of Finance is auditing. The set of rules defining the business' standard operations must be preserved and changes must be consciously approved and internally communicated to reduce the risk of misconduct. New protocols and policies defined temporarily to handle any emergent situation must also be documented. The utilization of decentralized governance mechanism applicable at three levels – decision making, rule management and auditing – have the potential to fulfill these requirements and yet be flexible enough to adapt to the dynamics characteristic of financial environments. Finally, in addition to self-explainability and auditability, the design aims at enabling the compliance system to act in accordance with the principles of transparency, fairness and accountability.

6. System Architecture and Deployment Scenarios

The architecture of the proposed approach consists of three main components: the centralized infrastructure, the various multi-agent systems employed for adaptive compliance and real-time risk prediction, and the interactions with external systems. Once the agent society is running and the first policies defined, the adaptive collaboration and multi-agent detection of anomalies should require limited human involvement. Security and privacy have an essential role, as these systems work with sensitive data in distributed environments. Communication protocols must ensure the confidentiality, integrity, and authenticity of data messages. Wherever possible, methods should be included to ensure privacy and confidentiality, allowing the exchange of public information with different degrees of confidentiality. Maintaining a trustful environment for real-time payments among partners with limited data sharing is essential. A compliance-by-design method considers the preservation of users' privacy.

The main objective is to offer an adaptive framework for multi-agent compliance policy definition, merging and refinement that takes into account actual transaction flows and schools of thought combining different detected anomalies. Example applications include banking transfers, payments among different partners, and securities operations. First, the creation of a new multi-agent adaptive compliance model manages the distribution, merging, and refinement of Central Bank Digital Currency (CBDC) transaction rules by using transaction patterns and specifics built by a Detection of Anomaly-based Anomaly Handling (DAAH!) model. A combined approach using two other independent multi-agent systems deployed for real-time CBDC transaction evaluation detects flow anomalies and supervises the use of known and dependencies-developing Risk Indicators. Then, the combination of the previous systems should help identify new or unknown Transaction Patterns that require investigations and related policy definition, merging, and improvement for the overall multi-agent society.





Fig 6: System architecture for a FLERSA service deployment scenario

6.1. Infrastructure, Interoperability, and Security

Research objectives and findings involve the design of a multi-agent infrastructure for adaptive multi-objective financial crime risk management, adaptive operational compliance, and real-time prediction and early warning of financial crime risks. Special focus is given to the multi-agent context of the approach. Adaptive operational compliance is attained through the distributed management and enforcement of compliance rules via a multi-agent architecture embedded into the operational ecosystem. A systemic multi-agent interaction mechanism allows situational awareness and real-time anomaly detection through cross-agent coordinated sensing.

Key concerns for the infrastructure supporting the multi-agent architecture include scalability, interoperability, security, and data privacy. Scalability is addressed in two ways. First, multi-agent systems by design support a distributed processing approach, which makes it possible to address many demands in parallel by adding agents without introducing unmanageable latencies in decision-making or processing. Second, privacy-preserving federated learning techniques mitigate data privacy risks while still allowing the construction of accurate and reliable predictive models by facilitating model training on multiple partially labeled data sources without pooling them. Interoperability is both a condition for the proposed architecture materialization and a design principle for agent instantiation. Data availability and quality are the primary justification for such consideration; agent behavior is modeled and implemented using supervised statistical learning techniques.



6.2. Privacy Preserving Methods and Compliance by Design

Compliance by design minimises privacy risks in the collection and storage of sensitive user data. Regulatory authorities should evaluate the possibility of establishing open-source and trusted repositories for transaction data to preserve sensitive user data while ensuring compliance enforcement by data reconciliation. Data privacy risks may also be addressed using advanced cryptographic and number-theoretical concepts that allow detecting sensitive user transactions without revealing the user's identity to the detection agents.


Interoperability of risk and compliance data models among the banking, payment and securities infrastructures provides opportunities for compliance enforcement without collecting sensitive user data. For example, identifying a cash outlier transaction originating from a user of an online payment service does not require linking this transaction with any personally

identifiable information about the real user of the payment service. A simple non-homomorphic number-theoretical equality check is sufficient to assess the property of being outlier, and it can be done while the user’s identity remains secret. The reconciliation of transaction data exchanged between the involved infrastructures further reduces required data storage while enabling effective anomaly detection.

6.3. Case Studies: Banking, Payments, and Securities


Adaptive financial crime compliance and real-time risk prediction are highly interconnected in today’s world. While a financial institution cannot operate in a regulatory vacuum, merely ensuring compliance and conducting risk assessments at a specific point in time without any feedback loops are no longer sufficient. The analysis of other participants involved in the same banking, payments, or securities ecosystem is just as important as monitoring the institution itself.

Case studies in these domains showcase the requirements for adaptive compliance and real-time risk prediction, including the necessary functionality and the decision-making surfaces for the multi-agent AI system. Such functionality and features can be developed individually for each area and integrated later. However, to prevent duplication of effort, it is advisable to either implement common functionality for all sectors first—such as explainability, auditing, governance, distributed agent rule management, and integrated anomaly detection—or develop multi-sector functionality, such as rule sets and detection models used jointly or impartially across sectors, while still maintaining flexible allocation for optimal performance.



7. Challenges, Risks, and Mitigation Strategies

In the context of multi-agent retroactive and non-retroactive financial crime compliance systems and real-time risk prediction, several key challenges and risks warrant attention. Scalability is a general concern for AI-based financial systems, which should also be designed to allow yield on investment in a reasonable timespan. Latency both in predictions and in compliance actions




can affect the expected usefulness of risk prediction and compliance systems. Such systems typically operate in trust-based environments, meaning that detection of trust-based compliance rule violations may prove especially challenging. In risk prediction, the quality of the training set is critical for reliable operation and maintenance; detection and supervision of potential mislabeling is thus key. Some types of financial data have high temporal and concept drift, and uncertainty in the labeled training set can have more dramatic effects than in other domains. The regulation of Finance is in near real-time, meaning controls are expected and required to occur at all times. The natural imbalance in the amount of anomalous vs non-anomalous behavior often requires an active learning support.

Just as prediction cannot determine the future with certainty, financial prediction cannot predict with certainty the risk of a transaction. Nevertheless, neither the regulatory requirement of clearly stated, automated, rule-based policies for surveillance, detection and prevention nor the fact that management, regulatory and societal pressures, combined with growing availability of data, systems and governance capabilities, make rational, explainable, risky decisions. Supporting decision-making does not eliminate the requirement for explaining and justifying the final decisions. The nature, extent and contextual effect of the support requires continuous evaluation against business and risk appetite objectives. Decision support is likely to so improve the label quality of, and therefore the training of, active learning systems that it may finally have a greater source of value than the decision support itself.

7.1. Scalability, Latency, and Trust

As in many machine-learning applications, it is critical to allow the components of the agent system to generalize well, i.e., to avoid overfitting. The samples used for training must therefore be representative of the problem domain, and care must be taken that the training and target domains remain sufficiently similar. In other words, concept drift such as seasonality effects must be taken into account, and when deploying supervised classifiers such as neural networks and random forests, additional mechanisms are required to detect significant changes in the data distribution or other symptoms of unseen conditions. If unanticipated events occur that the system does not fully understand, decision support must still be available. Pure anomaly-detection mechanisms are therefore essential. The smart agents collaborating in the adaptive architectures that can generate and maintain the compliance-monitoring rules are the same agents that provide critical contextual information to static classifiers or supervised regression methods and validate any predictions produced by other agents.

In any multi-agent system, each agent must interact with the others in real time. Thus, the volume of messages being processed can quickly grow to a level that creates bottlenecks. Optimization strategies, load balancing, and horizontal scaling must be considered to facilitate the efficient execution of decision-making processes. Latency along the interaction path must be kept small. Building trust is also essential. Agents implementing the same functionality (such as monitoring a specific rule in the compliance architecture) may initially be regarded with equal trust. As decision-making proceeds, however, one agent may prove to be more reliable or accurate than its peers. Trust values must gradually adjust to take this information into account. Any system capable of predicting reputation for gig-workers—based on that of the workers' previous clients—could serve as a model for a multimodal Multi-Agent AI system.




Area	Article Emphasis	Evaluation Criteria / Metrics	Main Risks / Constraints
Adaptive compliance effectiveness	Measure how well the MAS satisfies compliance obligations in dynamic environments	Response time, anomaly detection quality, operational fit, audit readiness	Rule complexity, latency, fragmented compliance regimes
Real-time risk prediction	Assess predictive quality of timely risk-scoring models	Precision, recall, F-score, repeated reassessment, space/time-based split testing	Mislabeling, class imbalance, unstable labels
Simulation environments	Use representative banking, payments, and securities scenarios	Benchmarking under realistic ecosystem conditions	Limited generalizability to all real-world settings
Infrastructure and interoperability	Ensure scalable and cross-system deployment	Parallel processing capacity, interoperability, security, privacy preservation	Bottlenecks, data-sharing barriers, integration cost
Explainability and governance	Support acceptance by regulators, law enforcement, and institutions	Narrative explanations, audit trails, transparent decision rationale	Low trust in opaque AI systems

Table 3: Evaluation, Deployment, and Key Challenges

7.2. Data Quality, Labeling, and Concept Drift


Timely risk scoring requires diverse external data related to customers, their transactions and activities, and the ecosystem in which these occur. Models that rely on such data are subject to changes in the underlying inputs. Unlike “data drift,” which refers to the quality of the data sources used to train a model, “label drift” refers to the degradation of the labeled instances over time due to changes in the underlying distribution of the class label. Because the detection of potentially fraudulent instances is much more challenging than that of normal instances, label drift can be a major problem. New labeled instances corresponding to true frauds may still be scarce for modeling and detecting rare events.

These issues highlight the necessity of continuous monitoring of the inputs and outputs of risk-scoring models. However, labels for fraud detection are subject to possible changes over time, as difficult-to-spot fraudulent events may later emerge as false negatives. Special attention must therefore be devoted to the validation of the labels used to train and predict, especially concerning the false negatives. Such validation should preferably be carried out in collaboration with domain experts, with whom the organization needs to maintain a close relationship.



7.3. Ethical and Legal Considerations

The solutions are primarily focused on adaptive compliance, using shift detection and retraining mechanisms for rules and risk prediction models. The practical use of such multi-agent systems in real-world environments, however, raises ethical and legal questions that need to be addressed. In distributed systems that allow individual agents to act relatively autonomously, problems of fairness and bias may occur, as different valid but incompatible system functionalities (e.g., privacy versus monitoring) could cause conflicts in the solutions proposed by the cooperating agents.



Most such negative effects are covered by the explainability and audit capabilities of the underlying supervised learning models (e.g., for risk prediction). A serious but critical aspect of these systems should be considered: agents may not only ask questions but also answer them. It should therefore not be taken for granted that all bank agents, for example, follow their own CSR principles. Also, the knowledge based on dubious data may be deliberately inaccurate. Although such behaviours are detected by monitors, they can lead to downstream implications that have not been analyzed. In situations where the action of a specific agent is considered harmful (e.g., sanction declines from a credit authority), a fine or disadvantage for this misbehaving agent may be legitimate because it does not coordinate with the behaviours of others or with the overall principle of being responsible.

8. Evaluation Framework and Experimental Methodologies

A comprehensive experimental methodology has been established to evaluate the performance of Multi-Agent AI (MAI) systems, addressing both adaptive compliance and real-time risk prediction capabilities. Evaluating the complete system in realistic settings is challenging due to the complexity of compliance requirements and the collaborative nature of agents. Thus, an evaluation roadmap with simulation environments and benchmarks has been designed to facilitate the systematic evaluation of individual primary capabilities, namely adaptive compliance and real-time risk prediction. The roadmap defines profiles for each analysis, establishing potential threats to the experiment's validity and limitations affecting the generalizability of results.


The insights gained from these experiments can mitigate the associated risks. An artificial financial ecosystem, consisting of multiple connected financial institutions engaged in various transactions, is established to evaluate the real-time prediction of financial crime risk. Agents in charge of risk scoring make predictions based on transaction and financial behaviours in real-time. To optimize the prediction pipeline, latency-averse and explainable scoring models are built and evaluated. The findings of this experimental approach can inform future model development.

8.1. Simulation Environments and Benchmarks

Dedicated simulation infrastructures emulate the environments of distinct financial services segments and measure compliance effectiveness and risk prediction ability through defined performance criteria. Building on the existing body of work, simulations of two representative segments are defined: banking with credit and liquidity risk prediction, payments with AML risk prediction, and securities with insider trading detection.

The banking use case involves transaction data from a European bank and considers credit and liquidity risks during the prevalent crisis caused by the COVID-19 pandemic. The AML use case focuses on transaction monitoring for a European payment service provider and predicts transaction-level risk for fraud and money laundering. The 证券市场用例管理层监控市场操纵，在线学习用于捕获潜在的概念漂移。 Simulations support open-source reproduction of ML models and allow the deployment of any AI model as a service for compliance effectiveness evaluation. The results of these simulations guide investigations into the multi-agent architecture, capabilities and method definitions for compliance by design, real-time collaboration, and explainability.


Types of White-Collar Crimes



8.2. Metrics for Compliance Effectiveness and Risk Prediction

The evaluation of adaptive compliance effectiveness is inherently complex and relies on multiple indicators linked to the specific multi-agent configuration and its operational context. Monitoring response times for compliance-related events across the entire financial ecosystem helps identify potential delays that may hinder compliance requirements from being fulfilled without disruption. The detection of endorsable data for risk prediction modeling is another relevant metric, as insufficient quantity and quality of such data will impair the overall risk scoring.

Nonetheless, a set of common metrics offers a useful decision-making foundation. For anomaly detection—integral to ensuring situational awareness and event sponsorship—the objective is to detect as many of the relevant anomalies as possible before



formal endorsement procedures resolve their classification, typically measured using precision, recall, and F-score. For the risk prediction module, the predictive quality of the support model is paramount, monitored via the so-called space and time-based split testing strategy, with modeling performance reassessed every time a new risk score is calculated.

8.3. Validation, Governance, and Reproducibility


Developing risk predictors for financial crime and compliance in real-time must take into consideration the multi-participant characteristic of economic ecosystems and the asymmetric distribution of data between the involved organisations. In this context, the audited simulation environment is composed of several interrelated participants including, for example, central banks, banks, fintech companies, crypto exchanges, and payment service companies, with the primary aim of fraud detection in the digital payment area. The risk prediction model can be trained on one ecosystem participant (e.g., payment service provider), but then used in the operational environment of other ecosystem entities (e.g., banks and/or other financial service providers). The auditing analysis considers the practical operationalisation of the banking sector's payment service ecology from a top-down perspective. Here, the formalised multi-agents framework develops simulated decision-support activities for the central bank acting as a supervisory authority for the payment service market.

The second type of auditing analysis involves the practical operationalisation of the payment service ecology from a bottom-up perspective. In this case, the digital payment service provider's decision-support activities include fraud detection and prevention using a supervised learning algorithm for reverse transfer scoring of ongoing digital transactions and decision recommendation for the operational environment. Validation of the fraud detection model is done with the aid of a false positive and false negative analysis. The third type of auditing analysis involves the practical operationalisation of the banking sector's payment service ecology from a top-down perspective. The formalised multi-agent framework develops simulated decision-support activities for the central bank acting as a supervisory authority for the payment service market. The emphasis of the proposed auditing framework and methodology is therefore on multi-collaboration activity and multi-decision support. The goal is to create a flexible and adaptive multi-agent-based risk detection, prevention, and prediction system that operates in a multi-set environment with a diversity of agents and octants.

9. Conclusion

The growing sophistication of financial criminals and their networks, combined with the relentless pace of technological evolution, are preventing existing compliance approaches from being effective. Despite continuous investments in compliance measures, criminals still seek success and regulatory authorities are concerned that precious resources are being drained to full-fill mere compliance requirements. Hence, the investment-to-incident ratio is rapidly deteriorating rather than improving as desired through technology adoption.

Multi-Agent AI systems can establish an adaptive, collaborative environment for compliance that leverages the capabilities of multiple competent agents acting as a community of intelligent entities. Individual agents undertake specialised roles with the autonomy, intelligence, and capability to adapt without manual intervention. Mission objectives are achieved through self-organisation and real-time collaboration, exploiting dynamic inter-agent relationships of mutual benefit. When deployed in financial crime compliance applications, such systems not only address compliance shortcomings, but extend beyond mere



compliance to enable real-time detection and assessment of risk throughout the financial ecosystem, thus going beyond compliance to risk mitigation.


9.1. Future Trends

The evolving ecosystems of financial crime compliance present a wealth of new opportunities for innovation. The introduction of new technologic methods and support of the academic community provide tools and models for improving traditional financial crime compliance processes. Multi-agent systems are an open, flexible, and powerful approach to achieve a significant enhancement of the financial ecosystem through cooperation between traditional financial institutions and new players. The separation into specialized roles, combined with common sense, allows for a clear separation of horizontal and vertical complexity in the design of these systems, allowing the development of smaller solutions that can be integrated into larger multi-agent systems.


The banking payments and securities ecosystems can benefit from new specialized operators that use technology and big data to deliver fresh and innovative services to customers and partners. The changes in social and regulation change the risks in the ecosystems and provide a unique opportunity to develop a real-time risk prediction service. The models can be used by different participants to check users and policies. Different anomalies can be detected and identified allowing timely alerts, and monitoring and governance audit information can be provided.

10. References

1. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, Article 100402.
2. Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M. E., Elshafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences*, 12(19), Article 9637.
3. Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms. *IEEE Access*, 10, 39700–39715.
4. 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>
5. Alsuwailam, A. A. S., Salem, E., & Saudagar, A. K. J. (2022). Performance of different machine learning algorithms in detecting financial fraud. *Computational Economics*.
6. 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.
7. Rocha-Salazar, J.-J., Segovia-Vargas, M.-J., & Camacho-Miñano, M.-M. (2021). Money laundering and terrorism financing detection using neural networks and an abnormality indicator. *Expert Systems with Applications*, 169, Article 114470.



8. 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.
9. Shahana, T., Lavanya, V., & Bhat, A. R. (2023). State of the art in financial statement fraud detection: A systematic review. *Technological Forecasting and Social Change*, 192, Article 122527.
10. Shou, M., Bao, X., & Yu, J. (2023). An optimal weighted machine learning model for detecting financial fraud. *Applied Economics Letters*, 30(4), 410–415.
11. Singh, A., Jain, A., & Biabie, S. E. (2022). Financial fraud detection approach based on firefly optimization algorithm and support vector machine. *Applied Computational Intelligence and Soft Computing*, 2022, Article 1468015.
12. 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)
13. Smith, Q.-J., & Valverde, R. (2021). A perceptron based neural network data analytics architecture for the detection of fraud in credit card transactions in financial legacy systems. *WSEAS Transactions on Systems and Control*, 16, 358–374.
14. Srokosz, M., Bobyk, A., Ksiezopolski, B., & Wydra, M. (2023). Machine-learning-based scoring system for antifraud CISIRTs in banking environment. *Electronics*, 12(1), Article 251.
15. Pocher, N., Zichichi, M., Merizzi, F., et al. (2023). Detecting anomalous cryptocurrency transactions: An AML/CFT application of machine learning-based forensics. *Electronic Markets*, 33, Article 37.
16. Hernandez Aros, L., Bustamante Molano, L. X., Gutierrez-Portela, F., Moreno Hernandez, J. J., & Rodríguez Barrero, M. S. (2024). Financial fraud detection through the application of machine learning techniques: A literature review. *Humanities and Social Sciences Communications*, 11, Article 1130.
17. 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.
18. Karangara, R., & Subbagari, S. (2024). Filtering and detection of anti money laundering with the aid of optimization-enabled machine learning techniques. *International Journal of Computer & Organization Trends*, 14(1), 1–11.
19. Ajagbe, S. A., Majola, S., & Mudali, P. (2025). Comparative analysis of machine learning algorithms for money laundering detection. *Discover Artificial Intelligence*, 5, Article 144.
20. Raja, M. R., Hosen, M. A., Kabir, M. F., Sultana, S., Ashraf, S. A., & Islam, R. (2025). Detecting and preventing money laundering using deep learning and graph analysis. *International Journal of Advanced Computer Science and Applications*, 16(6).
21. 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.
22. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2021). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*.
23. 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*.
24. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2020). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*.



25. 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>
26. Lucas, Y., Portier, P.-E., Laporte, L., He-Guelton, L., Caelen, O., Granitzer, M., & Calabretto, S. (2020). Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs. *Future Generation Computer Systems*.
27. Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2020). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*.
28. Gorle, R. (2026). *Adaptive by Design: Building Data Infrastructure That Rewrites Itself at Scale*. Apollo Publication.
29. Lebichot, B., Le Borgne, Y.-A., He-Guelton, L., Oblé, F., & Bontempi, G. (2021). Deep-learning domain adaptation techniques for credit cards fraud detection. *Proceedings of the International Joint Conference on Neural Networks*.
30. Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2021). Streaming active learning strategies for real-life credit card fraud detection: Assessment and visualization. *International Journal of Data Science and Analytics*.
31. 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*.
32. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2020). The application of data mining techniques in financial fraud detection: A classification framework and an academic review. *Decision Support Systems*.
33. 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.
34. Kshetri, N. (2021). Blockchain and cryptocurrencies: Economic and regulatory issues. *IT Professional*.
35. Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., & Leiserson, C. E. (2020). Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. *Proceedings of the ACM SIGKDD Workshop on Anomaly Detection in Finance*.
36. Eberhard, L., Papp, P., & Stütz, R. (2022). Graph-based machine learning for financial crime detection: A review of methods and applications. *Journal of Financial Crime*.
37. 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. <https://doi.org/10.1109/icsmartgrid70882.2026.11680767>
38. Chen, Z., Li, C., & Sun, W. (2021). Bitcoin price prediction using machine learning: An approach to sample dimension engineering. *Journal of Computational and Applied Mathematics*.
39. Bussiere, M., Fratzscher, M., & Müller, H. (2020). The financial crisis and the global financial system: Implications for financial surveillance. *Journal of International Money and Finance*.
40. 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. <https://doi.org/10.1109/icsmartgrid70882.2026.11680752>