ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

AI-Powered Real-Time Sanctions Screening for Cloud-Native Finance

Nareddy Abhireddy
Independent Researcher
nareddy.abhireddy.researcher@gmail.com

Abstract

Regulatory frameworks for sanctions compliance require organizations to screen customers, transactions, business partners, suppliers, and other entities continuously and without delay. Sanctioned-party lists are updated frequently and can contain hundreds of thousands of records. The true cost of non-compliance is not primarily financial fraud exposure but regulatory penalties—fines that can expose not just the organization, but also its directors and senior officers, to personal legal liability. Given these stakes, real-time sanctions screening has become a critical operational requirement.

For organizations built on cloud-native architectures, conventional approaches—staging data in a data warehouse before applying screening rules—are no longer viable. A cloud-native solution instead demands event-driven processing, embedding screening directly into real-time data flows. Moreover, strict latency and throughput constraints often exceed what purely deterministic rule-based systems can support, making AI and ML-based risk detection a valuable complement. This paper outlines practical guidelines for implementing real-time sanctions screening and examines the key architectural and technical considerations involved.

Keywords—Real-Time Sanctions Screening, Sanctions Compliance Systems, Event-Driven Compliance, Cloud-Native Compliance, Transaction Screening Analytics, AI in Sanctions Detection, ML-Based Risk Detection, Regulatory Risk Management, Watchlist Screening Systems, High-Throughput Screening, Low-Latency Compliance, Continuous Compliance Monitoring, Financial Crime Prevention, Compliance Automation, Dynamic Watchlist Updates, Streaming Data Compliance, Risk Scoring Models, Compliance Data Pipelines, Intelligent Screening Systems, Regulatory Technology Systems.

I. INTRODUCTION

Regulatory authorities and law enforcement agencies require financial institutions to screen customers and transactions against global sanctions lists, which are updated repeatedly. Delays in performing these checks expose institutions to substantial risks. Cloud-native architectures, designed to scale according to processing demand, should enable low-latency sanctions screening. At the same time, compliance functions in financial services are being transformed through automation and the use of advanced technologies, including artificial intelligence and machine learning.

Achieving a meaningful reduction in latency requires managing the key challenges associated with real-time sanctions screening: data quality, latency, throughput,

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

privacy, auditability, and cross-border compliance. Continuous sanctions screening seeks to provide repeatable, traceable decisions as events within the financial ecosystem evolve. Three central principles guide this approach: sanctions lists are continuously monitored for changes; data flow is incessant, with decisions being made at the earliest possible moment; and substantial effort is devoted to ensuring that data provenance can be established for every decision taken by the system.

II. Background and Motivation

Screening customers and transactions against sanctions lists is a regulatory requirement for banks and other financial institutions. These lists, administered by government agencies such as the Office of Foreign Assets Control, the United Nations Security Council, and the European Union, are lengthy and dynamic, with extensive explanatory text that requires human interpretation. However, many organisations perform the screening operation only sporadically, rather than in real time for every new transaction or customer relationship. Insufficient data quality, lengthy processing latencies, and high throughput demands often inhibit real-time screening. Institutions that do not screen in real time are still obliged to monitor for new sanctions data, but without automation or a structured framework such efforts risk becoming an afterthought.

A cloud-native architecture, based on microservices, event-driven design, and support for AI/ML technology, facilitates continuous sanctions screening; yet consideration of data quality, processing latencies, and alert throughput is essential for success. These requirements relate not only to system architecture or technology selection, but also to the supporting processes that drive the detection of sanctions-relating signals within the evolving data landscape. With this consideration in mind the three key challenges of Governance, Operations, and Data remain useful for framing the development of AI-driven screening technology.

Fig 1: AI In Sanctions Screening

A. Rationale and Key Challenges in Sanctions Screening

Governments impose sanctions on organizations or individuals to influence specific behaviours or conditions, such as the termination of repressive acts or unlawful territorial occupation. These sanctioning actions usually create lists of prohibited parties, across a wide range of use cases, which organizations need to check against in a real-time manner, as new sanctions can be introduced at any time. Any transaction involving a sanctioned party is illegal and exposes the organization conducting it to significant legal risk, especially with regards to regulatory compliance. Governments have moved to criminalize the failure to detect and report on transactions involving sanctioned parties, increasing the requirement to invest in sanctions screening so as to ensure that transactions can be completed with minimal delay.

Some organizations operate in regulated markets where sanctions screening is a legal requirement, others treat it as Good Practice because of the legal risks involved and the geographical spread of cross-border transactions. Yet, sanctions screening is a difficult risk to manage, for a number of reasons—data quality, latency, throughput, privacy, auditability and cross-border compliance. Legal and regulatory requirements differ from jurisdiction to jurisdiction leading to a patchwork of laws, regulations and Good Practice Principles and Guidelines. All organizations covering these legal obligations need to pay special attention to compliance with laws of other jurisdictions enforced in parallel with local laws, but from which they are exempt from external application.

Equation 1. Latency Equation

Start with event arrival time:

$$t_a$$

Processing completes at:

$$t_c$$

Therefore latency is:

$$L = t_c - t_a$$

For multiple events:

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

$$L_{avg} = \frac{L_1 + L_2 + \cdots + L_n}{n}$$

So:

$$L_{avg} = \frac{1}{n}\sum_{i=1}^{n} L_i$$

III. METHODOLOGY

The study adopts an applied research approach, applying existing but evolving technologies in the context of real-time detection of sanctions violations in cloud-native financial systems. These technologies include natural language processing and machine learning. The data used to support feature-engineering, machine learning, and risk-analysis capabilities are made available for external use. The research is supported by a custom data set created to focus on the mining of sanctions-monitoring signals from news articles. To assist in the rigorous assessment of ML models, an evaluation framework is proposed that is aligned with the Machine Learning Lifecycle defined by the Institute of Electrical and Electronics Engineers.

Data Preparation, Feature Engineering, and Model Selection Techniques. The detection of sanctions violations poses unique problems that need to be addressed to realize efficacy gains through AI/ML development. These approaches aim at minimizing false negatives, supporting advanced risk-monitoring solutions, and understanding how the risk landscape is evolving through time. Successful adoption requires careful attention to design and integration aspects, ensuring that the compliance function is not only effective but also efficient, allowing resources to be allocated to areas of greater priority and risk. The primary detection task is supported through data integration and feature engineering, allowing other modeling attempts to yield higher precision, enhance risk monitoring, and assess the supporting data qualities.

Table 1: Core Research Themes

Area	Article Focus	Measurable Variable
Latency	Real-time screening	Decision time
Throughput	High-volume event processing	Events per second
AI/ML	Risk detection	Risk score
Compliance	Traceability and auditability	Audit record completeness
Privacy	Masking and pseudonymization	Data exposure reduction

A. Essential Approaches for Optimizing Sanctions Screening

Scalable cloud-native environments, ideally founded on microservices architectures, support large-scale adoption of all types of technology by ensuring rapid and cost-effective deployment, flexibility, and future-proofing. Yet few cloud-native financial services have embraced artificial intelligence or machine learning for risk detection or mitigation. Three well-established techniques are essential for optimization of timesensitive sanctions screening within such environments: First, a centralized repository for sanctions signals should be maintained as closed as possible; second, a comprehensive list of features presenting credible evidence of present sanctions risk should be compiled; and third, the solution should have an identified owner (or owners) with sufficient seniority and oversight responsibility to ensure its proper, reliable, maintained, operated, and audited.

Sanctions signals, whose integration supports the world's most widely adopted real-time sanctions-checking technology, also correspond to sanctions alerts that propagate into business operations in violation of the Financial Action Task Force's immediate-reporting requirement. Combined, these signals imply an inherent lag of hours or even days in sanctions checks. Recent developments offer new possibilities for sanctions screening and sanctions detection: a recognised interest in real-time sanctions screening, together with the rapid, widespread deployment of cloud technologies and the clear adoption of AI-based solutions for risk detection beyond sanctions screening.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
ISSN: 3067-1140
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

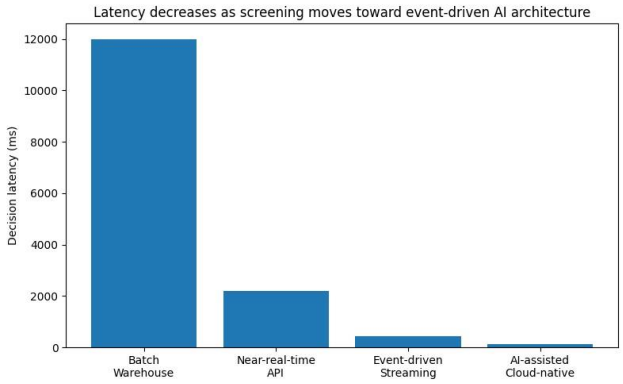

IV. Objective of the Study

Evolving technologies can support regulatory compliance in cloud-native architectures at scale, yet they are not universally deployed. Latency issues in data-stream processing hinder compliance with immediate sanctions listings. Timely compliance reduces financial institutions' risk exposure as new breaches emerge, but sanctions screening is often not performed in real time: data are batched, analysis deferred, or management reporting performed without immediate follow-up. In practical applications, below-acceptable data quality in screening and high false-positive levels during processing further hamper compliance.

The analysis aims to enhance regulatory compliance through scalable cloud-native solutions that use AI, ML, and advanced data-processing techniques to strengthen sanctions listing screening. Continuous, automated processes minimize reliance on manual intervention and fulfil ethical duties of care and transparency by formalizing reasoning for sensitive decision-making against regulatory lists. Reinforcement of the principles of AI/ML in advanced decision-management techniques for sanctions screening follows. The expectation is that controls and conditions will be established or available to ensure that routing, auditing, and justifying of sanctions decisions conform to best-practice fairness standards. The goal is to provide a repeatable, traceable, and dependable basis for privacy-driven sanctions detection.

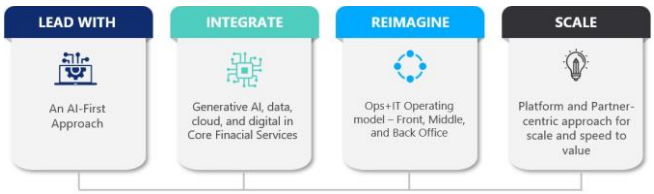

Fig 2: AI-Powered Financial Services

A. Enhancing Compliance through Advanced Technologies

Complying with sanctions such as those imposed by the United Nations, European Union, United States, and United Kingdom is a requirement for financial institutions, and cloud-native financial systems can leverage state-of-the-art technologies to improve the effectiveness and efficiency of these solutions. Technology can help reduce the cost of a compliance function, but cost is still an issue, especially with regard to detecting known indicators of a high-risk sanctions exposure. Timely, accurate, fair, and transparent responses to sanctions are essential to a robust and effective compliance function. Indeed, clients expect such responses, and management should be ready to explain how they are achieved.

The delicate matter of response latencies is often neglected, yet has important operational and reputational implications. Shorter latencies are of utmost importance for certain use cases, more so than for others, and yet they have received little attention in the literature. Continuous processing of data streams rather than discrete processing of batches allows reduction of latency, and positioning decisioning as close as possible to the data streams reduces latency even further. Latencies in high-frequency trading are on the order of nanoseconds, but much higher tolerances also exist: capital markets firms routinely close books on the day, often partially automating the process. Repeatability is essential for longer latencies, as transactions enter evergreen databases that are often subject to further enrichment before an eventual response.

Equation 2. Throughput Equation

Throughput means number of screened events per unit time.

If:

$$N = \text{number of events}$$

and:

$$T = \text{processing time}$$

then:

$$Q = \frac{N}{T}$$

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

If 22,000 events are screened in 1 second:

$$Q = \frac{22000}{1}\, Q = 22000 \text{ events/second}$$

V. Research Summary

The research focuses on sanctions screening within the context of compliance for cloud-native financial systems, where algorithms can optimize the timeliness and cost of this function. A principle of continuous sanctions screening underpins the exposition: the proposed systems should screen new data as it arrives and retain the decisions with sufficient contextual information to support an audit trail. Automation is essential to mitigate risks associated with sanctions lists, which can change rapidly and require prompt adaptation. Three key principles guide the research: repeatability of screening decisions to enable independent validation and a fully compliant audit trail; traceability that makes it possible to reconstruct and review a screening decision; and rapid adaptation of screening mechanisms as list entries are added, modified, or deleted.

In practical terms, these principles imply a focus on low-latency sanction screening decisioning; rigorous data governance to support screening at scale and across business lines; and an architecture that continuously monitors for new sanction signals while processing data through the screening pipelines. Continuous, real-time, and hitless sanctions screening is an aspiration of the financial-services industry. Several technology trends have combined to make this vision achievable for organizations operating cloud-native systems: the new-generation architectures provide the modularity and interoperability needed to ensure systems can be continuously refreshed without disrupting business operations; modern cloud computing approaches automatically scale processing resources to meet spikes in data volumes caused by product launches, regime crises, and other events; and a variety of artificial-intelligence and machine-learning techniques can be applied to optimize screening decisioning.

A. Key Principles of Continuous Sanctions Screening

Sanctions detection is a dynamic process that demands constant attention. However, the burden of exhaustion has overshadowed the principles of repeatable execution, auditability, and governance, crucial for any compliance-related activity. Continuous sanctions screening must evolve beyond the reactive model relied on by organizations at least at the top tier. The approach is akin to continuous fraud detection of payment instructions and counterparty risk in transactions.

Any mature architecture must support low-latency decisions, prompt and perspicuous data governance, and yearn for constant activity, like the canaries that signal danger to miners. Sanctions lists constantly evolve with additions, deletions, and metadata changes, creating a constant flow of updates. Continuous activity and speed indeed yield a clear advantage: lurking problems visible in historical execution can often be anticipated and avoided. In the context of continuous sanctions screening, these three important aspects stand out: the need for low-latency decisioning, the establishment of appropriate data provenance with minimal friction, and the realization and detection of sanctions signals on an event-driven basis to the extent possible.

VI. Principles of Real-Time Sanctions Screening

Maintaining the relevance and accuracy of sanctions screening in the evolving landscapes of risk, finance, and compliance across global jurisdictions constitutes a complex challenge in cloud-native implementations. Addressing temporal latency between transaction processing and sanctions list mutations is foundational to supporting ongoing relevance. Nevertheless, simply reducing compute latency to near-zero is not enough—actions taken must be both timely and correct based on the information available at the time of the decision. The principles underpinning continuous sanctions screening of flow-based data streams declare that it should be repeatable, should allow tracing of the conditions and decisions made, and should adapt rapidly to changes in the sanctions lists in use.

Sanctions checks that consume information with inherently low latency may almost be performed in real time if done on a calculation basis for each decision and with no knowledge of what checks have occurred previously. Timing of call execution or processing is one aspect of latency. Such decisions may be based on different data access conditions, with some implementations requiring specific copies of data sources to remain available for low-latency action. Real-time decisioning requires that any piece of information is available to routine systems at the moment it is needed, either through a direct copy from those systems or some reliable near-simultaneous message to sources that they may make a decision based on that information during processing.

ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

Fig 3: Agentic AI Use Cases in Banking

A. Key Strategies for Effective Sanctions Screening

Supporting continuous sanctions screening involves multiple strategies that affect system performance, efficiency, and relevance. The first key decision involves implementing screening with the lowest possible latency, promoting timely alerts and mitigating risks that arise from sanction matches after transactions have settled.

Any sanctions screening solution is dependent on the availability and integrity of the data sources feeding it. A second strategy therefore concerns data provenance, all aspects of data acquisition, enrichment, and delivery that together determine the reliability of risk signals. For real-time processing, a third consideration is a mechanism for continuously monitoring known and expected sanctions watching brief data, ensuring that screening is performed literally in real time—upon generation or receipt of an event, document, or activity—rather than at a later time of low activity volume.

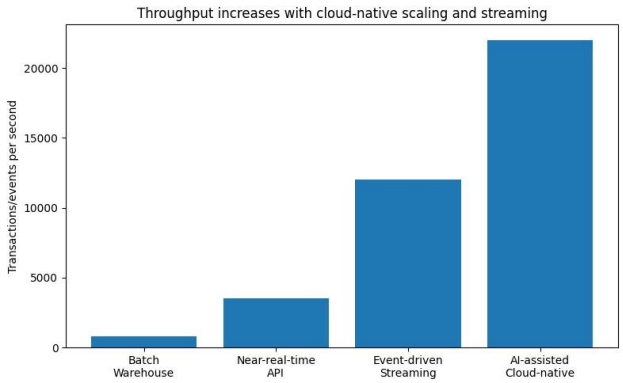

VII. Cloud-Native Architectures for Compliance

Throughout the previous sections, compliance was addressed in the context of an AI-enabled—partially or completely—sanctions screening capability but the focus on compliance was there. Any compliance-related capability is also an ideal candidate for a cloud-native implementation, enabling a continuously-compliant organisation in its deployment. The importance and novelty of a cloud-native compliance capability in a financial institution arise from three attributes: the real-time or near-real-time need of compliance with most regulations; the high volume of processing, which may induce occasional bursty behaviour; and the low tolerance for privacy violations, as significant volumes of private data from the clients are processed for compliance purposes. The cloud-native architecture—and thus, compliance capabilities—need to be set up, implemented, and deployed in adherence to an accepted posture and within a framework established by the organisation.

A cloud-native architecture comprises long-lived applications and containers. Long-lived applications comprise several—typically thousands—of components and thousands of continuously executing processes across multiple public clouds, private clouds, and on-premises environments. The composition and orchestration may be controlled by condensing large traffic flows into manageable burts; this is achieved through a powerful and seamless API-driven approach for flexible inter-service communication, driven by reusable service contracts covering the semantic interchange among services and versioning of the service communication interfaces in line with the governance processes within the organisation. A cloud-native architecture relies on an event-driven system, typically accompanied by a streaming architecture. This latter embraces the real-time information flow processing requirements of the system and enables the handling of bursty data traffic in a distributed way without overcommitting resources.

A. Microservices and API-Driven Integration

Cloud-native systems embrace a modular design that underpins remote-team operation and provides the means to concurrently develop, deliver, and maintain independent components. Microservices have emerged as the leading architecture within this approach. At a high level, a microservice is a self-contained unit that implements a cohesive set of features and meets well-defined service contracts. These contracts expose functionality through APIs,

ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

serving as the integration mechanism between microservices and external entities.

API contracts define the expected request and response message formats and support the evolution of microservice functionalities with controlled deployment risk. Changes to a contract could pose compatibility challenges for clients and service consumers, so versioning is used to allow users time to adapt to any modifications.

Real-time processing of high-velocity data streams often requires rapid, adaptive decisions. Event-driven architectures facilitate this by reacting to business events in the data and propagating the outcome to interested parties. The architecture must be geared to ingest and react to data—potentially from multiple sources—at volume and at speed.

B. Event-Driven Processing and Streaming

Cloud-native compliance systems encompass a rich ecosystem of specialized components working independently while sharing services and data. Event-driven operation and support for data streams are key characteristics that facilitate low-latency execution and the ability to respond quickly to bursts of high activity.

An event is a discrete, recordable occurrence, such as a market transaction, a change in account information, or the arrival of a new sanctions list. Events can also signify the passage of time, making it possible to trigger periodic actions such as monitoring for vulnerabilities in networked systems, bots, and web applications. Events may be taken from an event log or a message broker—collectively termed a message bus—that queues them for later processing.

Cloud-native compliance solutions can consume and produce events through an asynchronous messaging paradigm, ensuring loose coupling. The initial component of such a system ingests data streams from external sources, while subsequent stages process and produce events of interest, which are either stored for later use or fed to other processing stages in real time. Stream processing models the system as a series of consecutive, infinite windows over event streams. Each time a new element arrives at a window, all processing components in that window respond, allowing the system to behave as if executing in a continuous loop. By suitably shaping and ordering these windows, it is possible to support both the detection of immediate sanctions risk and a simple form of temporal reasoning about sanctions alerts.

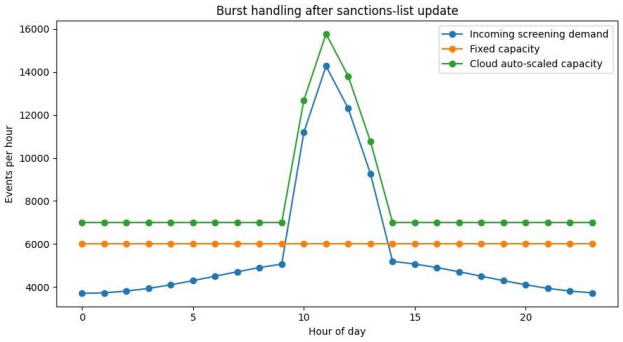

VIII. AI AND MACHINE LEARNING IN SANCTIONS DETECTION

Detecting sanctions-related risks requires a wide range of signals, and the right features must be engineered accordingly. Different risk categories signal distinct risks, and AI/ML techniques should ideally be incorporated into the screening process within an overall governance framework. Feature engineering and model training can be done in close collaboration with domain experts to improve detection quality, alleviate the risk of false positives, and operate with high levels of automation. Risk categorization can guide the search for wealth and income clues, hidden sanctions, and links to sanctioned entities. What the AI module needs, as in any supervised machine-learning exercise, is a well-structured labelled training dataset that supports the learning of patterns indicative of each risk.

Feature enrichment augments the original screening dataset with supplementary data derived from external data sources. Such features serve as additional indicators of sanctions or sanctions-related risk that may not be evident from the base data. Aggregate risk indicators roll up signal strengths across multiple features, domains, and edges. Training, evaluation, and deployment follow typical ML processes, with governance applied throughout to ensure that the models are valid, reliable, and free from bias. Bias is especially important in sanctions screening, as there are heightened social, ethical, and governance responsibilities to ensure fairness and convey the reasons behind the risk assessment. AI-based decisions should be interpretable, explainable, and auditable.

ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

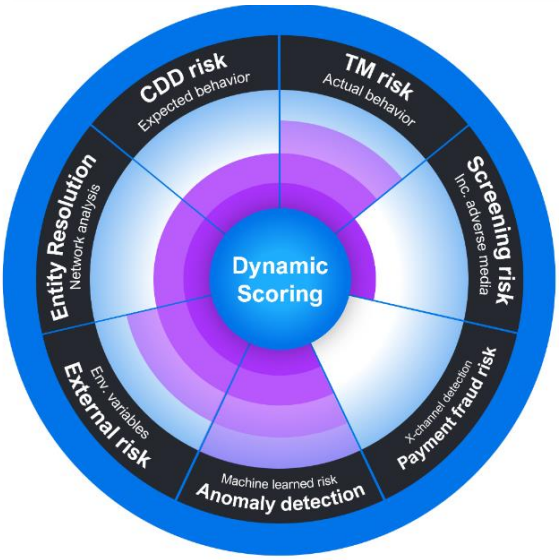

Fig 4: Dynamic Scoring

A. Feature Engineering for Sanctions Risk

Sanctions risk analysis draws from diverse sources to detect potential links between screened subjects and sanctioned individuals and entities. Control signals typically reside in legal data feeds that identify individuals and organizations subject to sanctions; threats arise from known or suspected affiliations with sanctioned organizations or connected actors. Official sanctions lists serve as direct indicators of affinity with sanctioned entities, while sectors subject to sanctions trigger exposure alerts. Supporting data enrich screening inputs and improve detection fidelity by reflecting patterns of FMIs engaged in risky transactions.

Feature engineering is crucial to comprehensively capturing the sanctions "signal." A streaming data architecture assembles signals continuously, exploring their practical effect on sanctions risk over time. Multiple layers provide redundancy and depth, with primary features assembled from public databases and alerts on sanctioned entities and known or suspected affiliations. Supporting and auxiliary data conditions build on aggregated labels derived from ICA clients and FMIs, with additional aggregation across organization networks.

Control factors are keyed to FMIs' transaction history. The sanctions risk score assigned to subjects and counterparties by an information provider offers explicit features for cross-border transaction screening. Aggregated transaction data examine volume characteristics, including counterparties engaged in cross-border transactions consistent over time. Signals of FMIs' links to Russian individuals and organizations and exposure to state ownership form other alert categories. Enhanced links gather data from other market participants.

Equation 3. Risk Score Equation

Each screening feature has a weight and a feature value.

$$R = w_1x_1 + w_2x_2 + w_3x_3 + \cdots + w_nx_n$$

Using summation:

$$R = \sum_{i=1}^{n} w_i\, x_i$$

Where:

Symbol	Meaning
R	Final sanctions risk score
w_i	Weight of feature
x_i	Feature value
n	Number of features

Example:

$$R = 0.30x_1 + 0.22x_2 + 0.14x_3 + 0.16x_4 + 0.10x_5 + 0.08x_6$$

B. Model Training, Evaluation, and Governance

Supervised AI models require carefully prepared training, test, development, and validation datasets. Decisions addressing these needs rely on earlier feature-engineering efforts, which calibrate the inputs presented to the models. These include designing combined feature signals that summarize the information derived from external risk-enrichment sources, the temporal horizon covered by their aggregation, and the frequency with which they are computed, sourced, and incorporated into the models. Once the feature-building approach is settled, the key risk-enrichment data layers can be identified.

Feature engineering for AI-driven sanctions screening involves building tailored features intended to improve coverage and accuracy. The success of a sanctions detection model depends strongly on these inputs. The goal is to extract

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

different kinds of information that explain why the entity holds the potential status being evaluated, e.g., is it being watched by an authority, are its country and/or industry being scrutinized, is its geography relevant for an RN detection given primary biases, etc. Exploring all this angle of information derived from external data sources is essential in order to allow machine learning algorithms, which are not naturally expert in addressing this problem, to discover the right patterns. Furthermore, big-data sources constitute a knowledge base that, if used correctly, can improve performance.

IX. Data Streams, Latency, and Throughput Requirements

Real-time sanctions screening imposes stringent demands on data streams, requiring low-latency processing to meet response-time objectives. In the UK, significant resources are reportedly devoted to sanctions compliance, making it essential to satisfy banks' monitoring needs efficiently. Cross-border detection of sanctioned entities can be supported through aggregate signalling. Addressing these factors helps ensure that applications can operate at the required throughput.

Global financial institutions typically maintain IDs aligned with data streams from sanctions-data providers. Recent incidents underscore the importance of data quality; therefore, two layers of integrity checking—one at the data supplier and the other at the user—facilitate prompt detection of compromised feeds. Secondary and tertiary feeds add redundancy and assist in creating sufficient data volume for temporal reasoning.

In real-time scenarios, the audit function examines a window of recent activity rather than verifying all transactions for potentially prohibited sanctions. During the bulk of any 24-hour window, the system needs decision-making capacity to cope with accumulated traffic, allowing a bursting mechanism to process activities more speedily as traffic subsides. When a sanctions list changes (through the addition of a new target, for example), a prompt check of recent activity is also required, preferably against a versioned and stable dataset.

A. Global Data Feeds and Source Reliability

Maintaining an up-to-date set of global sanctions lists that reflect the most current regulatory requirements is essential for any sanctions screening system. Ingesting real-time updates to the underlying data feeds is crucial but, by itself, does not guarantee sanctions signals are generated in a timely manner. Owners of real-time systems must also ensure these updates present low latency, sufficient throughput, and the requisite level of data quality. Other factors that support timely sanctions signal generation may also be less obvious, remote, or indirect, such as the ability to perform temporal reasoning over the messages in the data streams. Their implementation should therefore be considered carefully.

For the screening of events, actions, transactions, or customer risk profiles that are continuously monitored and that are not under management review, the focus is on short message latency and low throughput requirements. Yet data feeds also serve periods of burst processing, where the need for screening exceeds the local throughput capacity of the architecture. Sub-optimal characteristics of the feeds during those times are manifested as a high risk of packet loss, due to either the application-server network stack or the network router dropping packets. The consequences may be greater latency in the final business decisions, reactions out of business-time, or, in extreme cases, a failure to detect a sanctions screening signal entirely. Such failure may not always be evident, as detection might fall within the monitoring threshold set forth by the owner of periodic system auditing.

Table 2: Example Screening Metrics

Metric	Formula	Purpose
Accuracy	(TP + TN) / Total	Overall correctness
Precision	TP / (TP + FP)	How many alerts are true
Recall	TP / (TP + FN)	How many true risks are found
F1 Score	2PR / (P + R)	Balance of precision and recall
Latency	Finish time − Arrival time	Screening delay
Throughput	Events / Time	Processing capacity

B. Temporal Reasoning and Real-Time Decision Making

Sanctions compliance, carried out by financial institutions in accordance with applicable laws and regulations, often involves monitoring customers or counterparties in order to determine whether they are linked to designated persons or entities. The traditional approach for supporting these

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

activities relies on periodic screening of the pertinent data sources, compared against the various sanctions lists. However, the information underpinning the screening deteriorates over time while the risk linked to a given entity may change substantially for a multitude of reasons: changes in the entity's management, activities, shareholders, business partners, owners, or even new developments in geopolitical relations. Consequently, screening does not identify an increasing number of entities that these organizations should avoid. A real-time or continuous approach, where alerts are issued as soon as new information comes to light, is thus required.

Such strategies require the integration of different types of data streams that can combine breaking news with both structured and unstructured third-party information. The development of infrastructure that can support this level of temporal reasoning and real-time decision making, both on the data side and on the screening side, is therefore a key consideration for continuous sanctions screening.

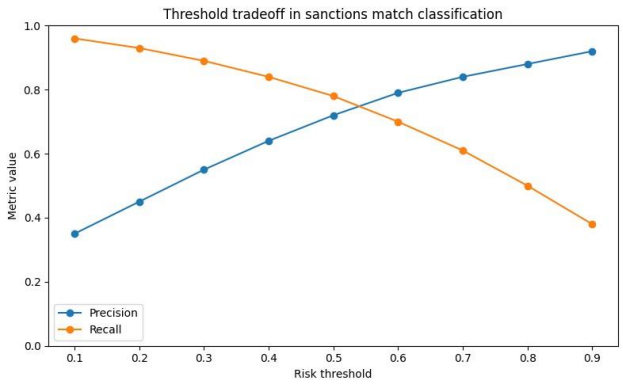

X. Compliance, Regulation, and Ethical Considerations

Compliance with laws and regulations exert critical forces on the design of real-time automated sanctions screening. System implementation of these requirements is necessary for legal sanction, and supportive features can improve effectiveness or mitigate risk. As such, ethics, governance, and compliance act both as constraints and guiding requirements. Together, they identify the stakeholders who should be consulted, the regulations adhered to, and the probable direction of investigation. The sanctions compliance obligations placed on financial institutions and other Covered Entities can shape the system. The need for proper classification, auditability, and record-keeping — in some jurisdictions, even the need to allow scrutiny of the decision-making process and outcomes by regulatory authorities — must all be respected.

A regulatory interface detailed enough for direct scrutiny can support auditability and compliance, reducing the need for separate audit processes and bespoke reports. Furthermore, compliance with laws focuses attention on areas where others might not, shaping oversight and guiding mitigation of risk associated with the system and its interactions with society. Fairness, bias, and explainability remain critical considerations in the machine-learning-everything era, and these elements can be captured by focusing on the communities that may be adversely affected by the outcomes of applying machine-learning technology. Government and regulators — such as bank boards, data protection authorities, and anti-corruption agencies — are important stakeholders and should be consulted in regard to fairness and bias.

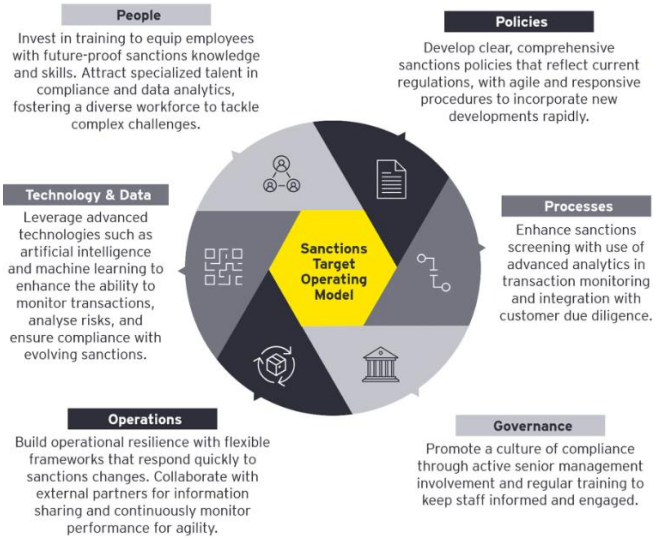

Fig 5: AI-powered sanctions compliance

A. Regulatory Interfaces and Reporting

Unlike the broader fraud and other risk domains, sanctions screening appears to be subject to a much narrower set of regulatory mandates and guidelines. At the Australian level, Dynamics 365 CE has been built to cover only the Australian Transaction Reports and Analysis Centre (AUSTRAC) reporting requirements. Likewise, Data and AI – Advanced Cloud, AI & ML for Financial Services has a specific focus on the Financial Crimes Enforcement Network (FinCEN) regulations. These elements are primarily

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

incorporated into the design to satisfy the relevant regulators, with minimal additional work involved to align with other stakeholders.

Bias mitigation, fairness, and ethical alignment play a particularly prominent role here, given the severe implications of erroneous risk indicators—indeed, the potential for sanctions screening and related security or governance functions to be used against entire socio-political groups has been highlighted by numerous commentators. Consequently, the ability of automated sanctions monitoring systems to not just provide accurate risk signals for business decision-making but also demonstrate that the systems are not inadvertently helping serve this purpose for the malicious actors remains paramount. Attention to these aspects stretches across multiple EU Global Data Protection Regulation (GDPR)-aligned requirements, especially those concerning the unprocessed outputs of associated processes. Consequently, these outputs are often fully audited and available to business decision makers, allowing them to mitigate risk and potential harm.

Equation 4. Alert Decision Rule

If risk score exceeds threshold:

$$R \geq \theta$$

then:

$$Alert = 1$$

If:

$$R < \theta$$

then:

$$Alert = 0$$

Complete rule:

$$Alert = \begin{cases} 1, & R \geq \theta \\ 0, & R < \theta \end{cases}$$

B. Bias Mitigation and Fairness

The decision-making process can introduce seemingly unjust discrimination against a stakeholder group, societal sub-group, or an individual. Bias may manifest due to the way data is labelled, the selection of training data, or the effect a feature has on the model scoring process. Common examples of bias might include: facial recognition algorithms that misidentify black people; algorithmic decision making for parole that recommends custodial sentences for black people more often than for white people with similar valuations for propensity to re-offend; or credit scoring algorithms that more often propose refusing credit for people in certain social demographic clusters than for people in broader society with similar attributes.

Those tasked with the product governance and risk management of machine learning models should validate that bias is either absent, or is documented, understood and properly managed through business as usual processes. For example when machine learning is used for loan scoring, departments with minority equity involvement should be aware of the probability of receiving an adverse action notice in order to consider if such probabilities should justify adjustment of the business policy (e.g. additional equity upfront), or the calculation of provision levels. With such information, business and model owner accountability is in place, and the process is better designed to satisfy their different sets of goals.

XI. Security, Privacy, and Data Sovereignty

Ensuring security, privacy, and data sovereignty throughout a sanctions screening architecture is paramount. Such considerations inform all design choices, affecting how the combination of microservices, API-based integration, and event-driven processing and data flows is secured from external threats; how data privacy and confidentiality are maintained; and how compliance with the relevant data protection regulations, including GDPR, CCPA, and HIPAA, is enforced.

Mapping the architecture onto the zero-trust principle requires defining who can access what and how. The confidentiality and integrity of secrets used in the architecture—for example, tokens for third-party data feeds—must be guaranteed, and access to those secrets must therefore be limited to those users and resources that absolutely need them at any point in time. Using a hardware security module for secrets management is advisable. To prevent unauthorized access to sensitive data, structured data

and especially the content of data flows must be masked or pseudonymized. A records system integrated into the screening architecture, and possibly separated from it, can act as a data safe, storing only essential information. Sensitivity and permitted use of the data can then be leveraged to

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

determine the level of aggregation needed for masking or pseudonymization.

A. Access Control and Secrets Management

These controls govern access to the system and its components. An authorization model built on the principle of least privilege is essential to mitigate damage from compromised credentials. For example, if an API must be invoked by specific services, the service account should be granted access to only those APIs instead of all APIs exposed by the cloud provider. A well-implemented secrets management process ensures that password-like secret values are prevented from being visible to any entity that doesn't need them, including developers. The Recommended Practices chapter provides additional guidelines.

Secrets management refers to storing and maintaining sensitive data such as passwords, encryption keys, OAuth tokens, API keys, and certificates. Systematic secrets management protects sensitive data from exposure during deployment (e.g., in published source code) and actual system operation. A best practice is to avoid hardcoding secrets or sensitive information into applications. Instead, secrets should be stored in a secure location such as secrets management services provided by cloud vendors (e.g., AWS Secrets Manager, Azure Key Vault) and retrieved securely by the application as needed. These services follow stringent security measures to protect secrets.

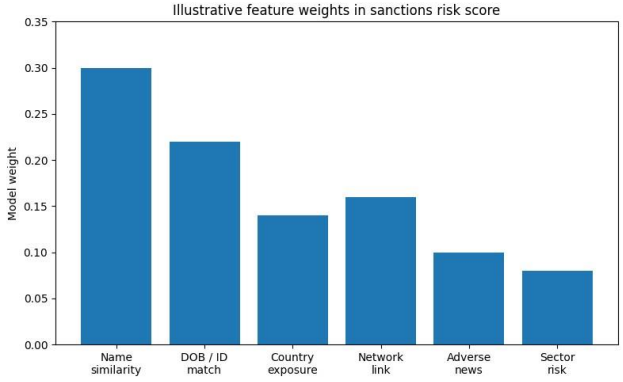

B. Data Masking and Pseudonymization

While privacy is of paramount concern, sanction screening can mandate access to sensitive details like addresses or phone numbers. Clearly, allowing unrestricted access can expose an organization and its clients to substantial risk, and thus a more pragmatic approach is warranted, one that balances risk exposure with operational agility. Sensitive field masking and pseudonymization is a common approach that supports data exfiltration while protecting client privacy. Integrating these techniques in a sanctions screening environment serves to reduce risk across the entire sanctions ecosystem.

Data masking replaces characters in a sensitive data field with masking characters, such as asterisks, preserving the field's size while maintaining a functional integrity during system integration. Pseudonymization replaces a sensitive field with correspondence to a value in a mapping table; the mapping is stored in a secure data vault, while the mapping table structure is stored as metadata. While pseudonymization can expose an organization to risk due to the need to maintain a vault, it can serve to mask sensitive details from analysts in a culture of open data sharing. From a sanctions perspective, masked foreign names, addresses or phone numbers can help preserve privacy for incoming clients without materially reducing the quality of screening analysis.

Equation 5. Precision Derivation

Precision measures how many predicted alerts are correct.

True positives:

$$TP$$

False positives:

$$FP$$

Total predicted positives:

$$TP + FP$$

Therefore:

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

Using graph values:

$$Precision = \frac{100}{100+430}\ Precision = \frac{100}{530}\ Precision = 0.1887$$

XII. RESULTS

Every process generates a large volume of information, which is critical for risk management and compliance functions. Testing and validation should therefore be

ISSN: 3067-1140
AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING
VOLUME: 01 ISSUE: 01
RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

thorough enough for a reasonable level of confidence to adopt the principles, technology, and approach. Specific requirements include accuracy, precision, and recall of the sanctions decision both on training and assessment data during a missing values experiment; data integrity check results during the monitoring phase; and support for the main international data privacy laws while using regulatory and KYC data.

The sanction screening process for a leading bank was subjected to a thorough validation and testing exercise. Results were collated in a format that addressed the validation requirements stipulated by the bank for similar approaches and technology. For each item being screened, a high-risk flag was set if: (i) the match score for any screening engine was higher than a pre-defined threshold level; or (ii) key details (e.g. name, date of birth) in the match list exactly matched keys in the KYC database. The results constituted a large imbalanced dataset; therefore, traditional accuracy was not an informative metric since the matches represented only a small percentage of the data throughout the study.

Testing the KYC preparation variations and missing values involved performing five-fold cross-validation repeated three times. The resulting two performance metrics clearly demonstrated the interest of the approach: high precision (mainly misclassifying non-matches) and also reasonably high recall (detecting a significant proportion of matches) for most KYC preparation variations and scenarios. Although the adoption of either of the initial versions of the method resulted in high true positives, the related precision was unacceptably low.

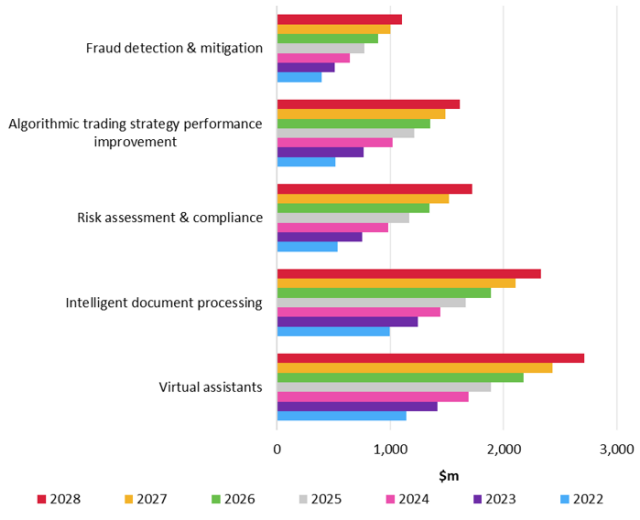

Fig 6: FS use cases in AI software spending, 2028

XIII. CONCLUSION

Real-time sanctions screening in cloud-native systems relies on three core principles: compliance must be auditable, responses traceable, and processes continuously repeatable. Satisfying these conditions enables regulators and regulated organizations to respectively review and explain any action taken within their respective purview. Real-time screening is not merely a speed requirement but mandates the continuous, low-latency assessment of events and transactions against available sanctions datasets. Continuous processing controls both the frequency of decisions and other aspects that govern suitability for real-time operation. Latency and decisioning tempo are thus distinct considerations. Although these demands are rarely defined as such by regulators, they must be satisfied for a sanctions-screening capability to operate as part of a compliance imperative.

Automation is a key enabler of low-latency processing. With the closed supply chain of riches, e-commerce, and marketing for nefarious activities offered by the dark web, the speed of exposure is accelerating, and that aspect far outweighs the risk of false alarms for transactors. The window between origination and validation is becoming so tight that smart criminals can slip through the gaps, especially in service-oriented dark-web trading. Speed of element data origin, data quality, and data life cycle — therefore speed of response — are paramount now. If real-time originate-validate capabilities are not in place, organizations will ultimately face extinction.

REFERENCES

1. Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Sudjianto, A. (2020). Detecting money laundering transactions with machine learning. Journal of Money Laundering Control, 23(1), 173–186.
2. Alarab, I., Prakoonwit, S., & Nacer, M. I. (2020). Comparative analysis using supervised learning methods for anti-money laundering in Bitcoin. Proceedings of the 2020 5th International Conference on Machine Learning Technologies, 60–66.
3. Yandamuri, U. S. (2022). Cloud-Based Data Integration Architectures for Scalable Enterprise Analytics. International Journal of Intelligent Systems and Applications in Engineering, 10, 472-483.
4. Kurum, E. (2020). RegTech solutions and AML compliance: What future for financial crime? Journal of Financial Crime, 30(3), 776–794.

ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

5. Lee, J. (2020). Access to finance for artificial intelligence regulation in the financial services industry. European Business Organization Law Review, 21(4), 731–757.
6. Canhoto, A. I. (2021). Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective. Journal of Business Research, 131, 441–452.
7. Mangala, N. (2021). CI/CD Pipeline Automation for Enterprise Data Artifacts Using Azure DevOps. Universal Journal of Business and Management, 1(1), 1-18.
8. 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, 100402.
9. Alkhalili, M., Qutqut, M. H., & Almasalha, F. (2021). Investigation of applying machine learning for watch-list filtering in anti-money laundering. IEEE Access, 9, 18481–18496.
10. Gupta, A., Dwivedi, D. N., & Jain, A. (2021). Threshold fine-tuning of money laundering scenarios through multi-dimensional optimization techniques. Journal of Money Laundering Control, 25(1), 72–78.
11. Loganathan, R. (2021). Integrated Risk and Compliance Frameworks for Global Data Center Operations: A Governance-Centric Approach. Universal Journal of Computer Sciences and Communications, 1(1), 1-26.
12. Gupta, A., Dwivedi, D. N., Shah, J., & Jain, A. (2021). Data quality issues leading to sub optimal machine learning for money laundering models. Journal of Money Laundering Control, 25(3), 551–555.
13. Domashova, J. V., & Mikhailina, N. (2021). Usage of machine learning methods for early detection of money laundering schemes. Procedia Computer Science, 190, 184–192.
14. Singh, C., Zhao, L., Lin, W., & Ye, Z. (2022). Can machine learning, as a RegTech compliance tool, lighten the regulatory burden for charitable organisations in the United Kingdom? Journal of Financial Crime, 29(1), 45–61.
15. Mattaparthi, R. (2021). Unified Data Lineage and Quality Governance Framework for Multi-Source Sensor Streams in Heavy-Duty Powertrain Manufacturing. Online Journal of Mechanical Engineering, 1(1), 1-15.
16. Kumar, A., Das, S., Tyagi, V., Shaw, R. N., & Ghosh, A. (2021). Analysis of classifier algorithms to detect anti-money laundering. In Computationally intelligent systems and their applications (pp. 143–152). Springer.
17. Koprivec, F., Kržmanc, G., Škrjanc, M., Kenda, K., & Novak, E. (2022). Screening tool for anti-money laundering supervision. In J. Soldatos & D. Kyriazis (Eds.), Big data and artificial intelligence in digital finance (pp. 233–251). Springer.
18. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. World Journal of Clinical Medicine Research, 1(1), 1-18.
19. Pei, Y., Lyu, F., Van Ipenburg, W., & Pechenizkiy, M. (2020). Subgraph anomaly detection in financial transaction networks. Proceedings of the First ACM International Conference on AI in Finance, 1–8.
20. Xia, P., Ni, Z., Xiao, H., Zhu, X., & Peng, P. (2022). A novel spatiotemporal prediction approach based on graph convolution neural networks and long short-term memory for money laundering fraud. Arabian Journal for Science and Engineering, 47(2), 1921–1937.
21. Mangalampalli, B. M. (2022). Automated Invoice Validation Systems Using Advanced SQL Analytics in Healthcare Insurance. Front Health Inform, 11.
22. Day, M.-Y. (2022). Artificial intelligence for knowledge graphs of cryptocurrency anti-money laundering in fintech. Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 439–446.
23. Stojanović, B., & Božić, J. (2022). Robust financial fraud alerting system based in the cloud environment. Sensors, 22(23), 9461.
24. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. Universal Journal of Business and Management, 1(1), 1-13.
25. Papantoniou, A. A. (2022). Regtech: Steering the regulatory spaceship in the right direction? Journal of Banking and Financial Technology, 6, 1–16.
26. Chao, X., Ran, Q., Chen, J., Li, T., Qian, Q., & Ergu, D. (2022). Regulatory technology (Reg-Tech) in financial stability supervision: Taxonomy, key methods, applications and future directions. International Review of Financial Analysis, 80, 102023.

ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 02 REVISED: OCTOBER 19 ACCEPTED: NOVEMBER 04 PUBLISHED: NOVEMBER 12

27. Reddy, V. A. R. (2021). Challenges in Standardizing Member Eligibility Data Across Multi-Payer Healthcare Ecosystems. International Journal of Medical Toxicology and Legal Medicine, 24(3), 1-19.
28. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. Expert Systems with Applications, 193, 116429.
29. Kumar, S., Ahmed, R., Bharany, S., Shuaib, M., Ahmad, T., Tag Eldin, E., Ur Rehman, A., & Shafiq, M. (2022). Exploitation of machine learning algorithms for detecting financial crimes based on customers' behavior. Sustainability, 14(21), 13875.
30. Moreira, M. Â. L., Rocha, C. de S., Silva, D. F. de L., Castro, M. A. P. de, Costa, I. P. de A., Gomes, C. F. S., & Santos, M. dos. (2022). Exploratory analysis and implementation of machine learning techniques for predictive assessment of fraud in banking systems. Procedia Computer Science, 214, 117–124.
31. Mangala, N. (2021). Optimizing Large-Scale ETL Pipelines Using Medallion Architecture on Azure Data Lake. Journal of Artificial Intelligence and Big Data, 1(1), 1-20.
32. Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., Elhassan, T., Elshafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. Applied Sciences, 12(19), 9637.
33. Błaszczyński, J., de Almeida Filho, A. T., Matuszyk, A., Szeląg, M., & Słowiński, R. (2021). Auto loan fraud detection using dominance-based rough set approach versus machine learning methods. Expert Systems with Applications, 163, 113740.
34. Rtayli, N., & Enneya, N. (2020). Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization. Journal of Information Security and Applications, 55, 102596.