

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



An Event-Driven Architecture for AI-Powered Financial Crime Surveillance

Bhasker Katta

Independent Researcher
kattaabhasker@gmail.com

Abstract

Financial crime detection and prevention remains a persistent challenge in the fight against criminal activity. The rapid growth in transaction volumes, combined with technological advances such as cloud computing, big data, the Internet of Things (IoT), artificial intelligence, and blockchain, has expanded the scale and complexity of complex event-processing (CEP) systems used to monitor financial activity. This same growth has also created new opportunities for malicious actors to exploit. Traditional surveillance approaches rely heavily on static, rule-based detection, which often fails to catch financial crimes that evolve over time — through changing actor behaviours, shifting patterns, and emerging threat landscapes that static rules were never designed to detect.

To address this gap, recent research has proposed event-driven regulatory technology (RegTech): a compliance paradigm that enables dynamic, real-time regulatory monitoring across complex event-processing systems. Building on this concept, we present FinCompliance 2.0, an intelligent compliance architecture designed to automate the detection of financial crimes with minimal human intervention. The architecture combines three core components: an event-processing layer, a flexible rule-based compliance engine, and intelligent detection mechanisms capable of adapting to evolving patterns of misuse. We evaluate the detection mechanisms against several benchmark datasets from the financial domain, and the results show performance comparable to state-of-the-art algorithms.

Keywords: Financial Crime Detection, Event-Driven RegTech, FinCompliance 2.0, Intelligent Compliance Systems, Complex Event Processing, Transaction Monitoring Systems, AI in Financial Compliance, Fraud Detection Analytics, Dynamic Compliance Monitoring, Regulatory Technology Systems, Real-Time Compliance Analytics, Financial Risk Detection, Adaptive Compliance Rules, Automated Crime Detection, Data Stream Processing, Behavioral Risk Analytics, Cloud-Based Compliance, Big Data in Finance, IoT-Driven Compliance, Predictive Financial Security.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



1. Introduction

For decades, the financial services sector has been heavily regulated, spawning a diverse range of specialized companies—RegTech, or regulatory technology—focused on compliance. Regulatory compliance, however, is often seen as a burden for regular businesses. The Covid-19 pandemic exacerbated this burden by bringing about rapid and unpredictable changes in customer behavior, business processes, and transactions. Existing compliance systems, often conceived years or even decades ago, can now no longer keep pace. Financial institutions, especially banks, hold vast quantities of often-untapped data, and it is now accepted in many domains that money laundering and fraud can be mitigated only through effective data analysis. Yet compliance requires not only data but also sophisticated event-driven processing of that data.

Event-Driven Regulatory Technology takes a prevailing compliance need, purpose, or event as a starting point and generates increased regulatory compliance through intelligent detection on data streams. Applying this approach to financial crime detection led to the design of FinCompliance 2.0, a generalized architecture for compliance with financial crime regulations. Bursting with untapped data, financial institutions can address compliance in new, more-effective ways. New intelligent-detection mechanisms such as anomaly detection, pattern recognition, or graph-based analysis allow examination not only of the “normal” level of compliance but also of changing patterns in compliance and the detection of sudden, latent, or long-term changes that disappear from conventional reporting.

2. Background and Motivation

The complexity and dynamics of modern business processes that span multiple domains make it difficult for compliance to rely on a predefined set of rules. A more suitable alternative is for compliance to shift from a reactive to a predictive paradigm where fraud and crime detection follow an event-driven approach. This is similar to the way an insurance company models risk based on a large number of events and transactions, detecting criminal behavior before it occurs. Event-driven regulatory technology (RegTech) caters for such fraud detection models by enabling the design and execution of adaptive compliance rules that evolve according to the metaknowledge available for an organisation at a certain point in time, and that detect events, events within environments, and complex events across several domains when certain conditions are satisfied.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



Research over the past ten years has shown that a comprehensive event-driven RegTech framework can assist banks and similar institutions to move compliance to a more intelligent detection level—beyond simple and metadata-driven detection toward anomaly, pattern, graph and network-based detection. However, these theoretical foundations only cover a subset of the requirements for a comprehensive framework capable of supporting intelligent detection in a real-world application. More specifically, the initial architecture lacked a dedicated event-processing layer that could take the disparate event streams generated by multiple monitoring systems and process them into a reasoned event stream that allowed both temporal reasoning and the detection of complex events.

2.1. Rationale and Impetus for Research

Although financial and cybercrime detection systems benefit from decades of research and extensive solutions based on various technologies (data mining, machine learning, artificial intelligence, and other available technologies), regulators and policy experts call for enhancing detection capabilities. However, the implementation of yet another datadriven detection method is often misleading. Exploring the limitations of data-driven approaches leads to active, event-driven regulatory technology (RegTech) systems that perceive, learn, and act. Eventdriven RegTech allows for intelligent detection capabilities and adopts anomaly detection, network, and pattern recognition approaches. FinCompliance 2.0 is an event-driven RegTech framework for intelligent financial crime detection that extends the research of intelligent detection to address effective compliance in a RegTech context.

Despite the successes, a persistent gap between the capabilities of detection systems and the growing sophistication of criminal actors remains. For example, some financial institutions process transactions with well-known secessionist or terrorist-related parties without supervision, while others fail during routine checks. Detection systems are often criticized not for false positives but for the number of undetected suspicious cases. The solutions are often developed in isolated silos of financial crime and cybercrime. Such large numbers of undetected cases are especially detrimental to the success of any financial crime detection project, as they undermine the purpose of detection that serves as an initialization mechanism for subsequent risk monitoring.



Fig 1: Benefits of Regtech in Financial Crime Compliance

3. Methodology

The methodology is structured into five major components. First, a summary of the fundamental concepts of event-driven regulatory technology is provided, together with an overview of architectonic layers of an event-processing system and key enablers for efficient supervision. Subsequent sections delve into the three main challenges of RegTech—data governance, intelligent detection mechanisms, and architecture—before placing focus on the two representative processing paradigms, covering data lineage, privacy-preserving techniques, and validation aspects.

The foundational aspects of event-driven regulatory technology serve to inform an architecture for intelligent financial-crime detection, FinCompliance 2.0. This system is viewed as an event-processing application that receives an event stream from source systems, offering intelligent detection capabilities and returning detected compliance events to those systems. At its core, an event-processing layer applies complex-event-processing principles, and its monitoring of temporal event patterns is complemented by a knowledge base containing compliance rules, semantic models, and trustworthy data. Surveillance against known behaviors is supported by an anomaly-detection mechanism, machine-learning methods, and risks associated with network or graph connectivity.

Equation A. Precision, Recall, and F1-score

The supervised detection evaluation using **accuracy** and **F1-score**, and later also refers to **ROC/AUC**.

1. Precision

Start from the confusion matrix idea:

- *TP*: suspicious case correctly detected
- *FP*: normal case wrongly flagged

By definition,

$$\text{Precision} = \frac{\text{correct positive predictions}}{\text{all positive predictions}}$$

So,

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

2. Recall

Also called sensitivity or true positive rate.



$$\text{Recall} = \frac{\text{correct positive predictions}}{\text{all actual positives}}$$

Hence,

$$R = \frac{TP}{TP + FN}$$

where FN is a suspicious case missed by the system.

3. F1-score

F1 is the harmonic mean of precision and recall.

Start with the harmonic mean formula for two variables:

$$H = \frac{2ab}{a + b}$$

Set $a = P$ and $b = R$:

$$F_1 = \frac{2PR}{P + R}$$

Now substitute $P = \frac{TP}{TP + FP}$ and $R = \frac{TP}{TP + FN}$:

$$F_1 = \frac{2 \cdot \frac{TP}{TP + FP} \cdot \frac{TP}{TP + FN}}{\frac{TP}{TP + FP} + \frac{TP}{TP + FN}}$$

Take the numerator first:

$$2 \cdot \frac{TP}{TP + FP} \cdot \frac{TP}{TP + FN} = \frac{2TP^2}{(TP + FP)(TP + FN)}$$

Now the denominator:

$$\frac{TP}{TP+FP} + \frac{TP}{TP+FN} = \frac{TP(TP+FN)+TP(TP+FP)}{(TP+FP)(TP+FN)} = \frac{TP[(TP+FN)+(TP+FP)]}{(TP+FP)(TP+FN)} = \frac{TP(2TP+FP+FN)}{(TP+FP)(TP+FN)}$$

Now divide numerator by denominator:

$$F_1 = \frac{\frac{2TP^2}{(TP+FP)(TP+FN)}}{\frac{TP(2TP+FP+FN)}{(TP+FP)(TP+FN)}}$$

Cancel the common denominator:

$$F_1 = \frac{2TP^2}{TP(2TP+FP+FN)}$$

Cancel one TP :

$$F_1 = \frac{2TP}{2TP+FP+FN}$$

So the final compact form is:


$$F_1 = \frac{2TP}{2TP+FP+FN}$$

3.1. Fundamental Concepts of Event-Driven Regulatory Technology

The primary innovation of the work is an updated version of the FinCompliance framework, which implements the core foundations of event-driven Regulatory Technology. In particular, it comprises an architecture that supports event streams and complex event processing, thus allowing compliance checks to be performed at all times on all events from various sources. The complex events detected in real time can trigger automatic and manual analytic inquiries that lead to powerful decisions. Unexpected events deviating from known flow patterns can also be detected using a combination of anomaly detection techniques and general-purpose simple rules.



Various systems, frameworks, and methods for Rule-Based Detection Mechanisms in Regulatory Technology have been proposed in recent years, yet very few solutions support the online detection of all forms of financial crimes by considering every business event from different data sources in compliance checks at all times. The only learning-enabled component is the automatic configuration of rules targeting the detection of attacks, while all other components also accommodate different types of training data used in supervised or unsupervised contexts. The idea is to go beyond a single detection mechanism based only on predefined rules, goals, or aims, by introducing a multi-pronged approach to achieve improved operational capabilities.



4. Objective of the Study

Aspect-based sentiment analysis is a special case of fine-grained sentiment analysis, comprising three principal tasks: aspects extraction, sentiment classification, and aspect-sentiment association. Aspect-based sentiment analysis focuses on the sentiments towards specific aspects or attributes of products, services, or entities. The aim of this relatively new research topic is to determine sentiment polarity—expressed as positive, negative, neutral, or one of the other assigned sentiments—both at the document level and at the aspect level.

Two recent and comprehensive reviews enabled the identification of several literature sublines. Specifically, Hu and Liu presented the pioneering approach, still ruling in recent years, that implemented both aspects extraction and sentiment classification based on knowledge bases and lexicons and thus classified the analysis as a top-down approach. Zhang et al. pointed out that the use of deep learning techniques to perform all three aspect-based sentiment analysis tasks at once has gained growing attention.

Table 1 — Main architecture extracted

Layer	What the article says it does	Main inputs	Main outputs
Event Processing Layer	Filters, enriches, transforms events; supports temporal reasoning and CEP	Transactions, customer data, account events, external alerts	Prepared event streams, correlated events
Compliance Rules & Knowledge Base	Stores compliance rules, semantic models, historical knowledge, governance logic	Regulatory rules, expert knowledge, historical patterns	Rule matches, semantic constraints, compliance triggers
Intelligent Detection Mechanisms	Runs anomaly detection, pattern recognition, network and graph analysis	Prepared streams + rules + contextual entities	Risk scores, suspicious-event flags, grouped alerts
Evaluation / Validation	Uses benchmark datasets and metrics such as ROC/AUC and F1	Model outputs and labeled benchmarks	Detection-quality assessment

4.1. Goals and Aims of the Research

The aim of this research is to create a next-generation event-driven regulatory technology framework capable of enabling intelligent violation detection for regulatory and supervisory purposes in a privacy-preserving manner. The underlying research objective is to design and implement an extensible RegTech platform called FinCompliance 2.0 that allows for the continuous detection of a wide range of financial malpractices in a timely fashion. The construction of such an extensible platform represents an initial step towards realizing the ultimate vision of general and intelligent RegTech by combining multiple appropriate, state-of-the-art techniques under an event-driven paradigm that supports diverse data-processing approaches. Data management constitutes another core aspect of FinCompliance 2.0. It features dedicated event-stream management mechanisms with built-in privacy-preserving techniques as well as careful consideration of data lineage.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



Financial-Service-Compliance 2.0 Detection of Malpractice-in-Network-based-complex-Event-FinCrime-DetectionsystemEvent-Sample-LeakApplication-Genericevent-Detectionopen-data-Event-Sample-LeakA-Benchmark Data Set for Money-Laundering Event Detection in Financial Networks-FinCompliance-2.0-A Multi-Literal-Classifer-ML-A-Based-Detection-Framework-for-Financial-Crime-Detection-Model-Driven-Principles-for-CrossDomain-Temporal-Evolution-Detection-in-SafetySupplementary-Sec-PR-Cone-Investigation-Presentation-TRDetecting-Moonlighting-and-Contract-Similarity-Resemblance-SchemesBased-on-MultiDatabase-Parallel-MineDetecting-E-Miner-Retrieving-Mobile-DangisetAl-Complicated-Illegal-Event-Detection-A-Specification-in-Contextual-based-Event-Strucure-Model-GraphConstitutive-measure-Sources-MedaLink-A-Matching-and-Detection-Framework-for-Online-MoveTracking-at-Micronecond-ElapsedTimeLink-Analysis-The-Trusto-Scandals-of-Walmart-and-MoneyGram-After-Purposely-Vertical-AnnexSupply-ChainCausation-Graph-Model-of-Digital-Transformation-Guided-Hetrogenous-Event-Driven-Event-AlignmmentMulti-Media-Corpora-Sentiment-Fusion-A-New-Approach-for-Multi-Media-MultiNarrative-Comparative-SentimentDetection-and-Comparative-Sentiment-Fusion-Support-A-Multitask-Semantic-Similarity-Detection-Net-Support-A-Context-Aware-Intelligent-Systems-Matching-Event-Eye-Making-Sure-The-Same-WaysSupported-On-The-Hidden-SecurityCAIA-Support-Detecting-Unsafety-API-Calling-Empirical-Evidence-from-MultiDomain-Crisis-Correspondingscome-A-New-Multi-Methylated-Single-Cell-Support-A-Multidimensionsupport-Damage-Detection-Remote-Sensing-Image-Detection-DamDetection-SmartPosture-Authentication-System-Using-Smartphone-Sensor-Data-Event-Derived-Total-OrolifinDecomposing-Networkx-XZ-Misspecial-Detection-Fault-Diagnosis-for-Event-SubgraphExtraction-A-Domains-Specification-A-social-activity-polarity-recognition-method-based-ondynamic-explicit-and-implicit-event-microblog-informationrecognition-Multi-Material-Perturbation-CrossViewTemperature-Dampening-Free-DetectionDeep-CrossDomain-MetaEvent-Detection-A-New-Paradigm-for-InterOutbreak-FinancialFraudDetection-Based-on-Dynamic-MultiGraph-Matchingwith-Applied-LearningAnalysis-of-word-embeddings-for-predictive-tweet-following-behaviorAn-event-structure-Based-Mixed-Method-of-KnowledgeAssociated-Event-Mining-Event-Sampling-A-CornerVisited-Based-Personalized-Destination-Event-DensitySurface-Event-Virtual-LearningDetailed-Control-and-Event-Sequencing-in-Water-Saving-Irrigation-DevicesDetection-WondemandChanges-of-Anorganic-Framework-and-Metal-Cations-During-Synthesis-andDrying-of-a-Pilar-12-MCM-41-Type-Material-A-68-GyrateCompletion-Network-and-PracticalRecommendations-for-DamageMinimization-Event-Detection-The-WSE-Blog-Net-Fiber-Detecting-and-Tracking-Moving-Events-in-RealTimeVideo-Cycle-Detection-Based-on-Even-TreeDecompositionof-

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



TimingDependent-Guided-MultiType-Event-Similarity-Discipline-A-Blend-
WaterCADConcentration-Event-Detection-A-Distributed-Algorithm-to-NearInstance-Optimal-
Detection-WithoutComplete-GraphSmin2-Eventtransparent-coverability-using-eventstructure-
Based-Logic-EventDetection-from-Traffic-Data-A-mobiledeviceGWG-Based-Distributed-event-
and-traffic-objectConformance-Check-Using-Event-RelationLogicalReasoning-on-Biological-
EventBased-Multi-EventType-EventDetectionof-Sensitive-Event-In-LargeScaleMmapping-
Multiview-AirPollutants-ClimatographyDatase3-Mapping-A-Hierarchical-Concept-EventFor-
EventTypeDetection-FusionShift-Net-A-3DVintage-Matching-Network-for3D-Economic-Event-
Digitalization-EventVertical-Deviation-GovernmentING-GovernanceAssigning-HumanBlind-
Turing-Machine-PreservingANew-Method-for-ColorSnapshotDigitalWatermarkingusing-Co-
Event1Phase-Event-Detection-AccurateLocation-Based-SocialNetwork-
EventRecommendationLeftEOFDetecting-Inherited-Vulnerabilities-using-DifferentialEDG-
from-MultivariateDyt-Event-WAP-Dcl-Timelocal-alther-a-multiple-sampling-basedChanel-
Event-Samplingan-EventDrivenApproach-to-DisasterAssessment-EventDetection-in-
OpenStreetMap-DataGain-New-Highlight-for-AgriNutrientTesting-TcantantiNode-
RecommendationBased-on-DominSigmaConceptin-DynamicUndirected-
SimpleGraphMiningCrossDomain1Datad-EventTypeReduction-A-
StackBasedApproachClustering-MultiSource-DrylandInformationinPolarAER-Detection-of-
HumanSecurityStoe-A-Deepcomposite-MultitechnologyMultiOrderDetection-based-
onsharingTypeplaceNameA-Rankconfidence-based-MultiMaterialExpandingRoutingLSTMD-R-
Detection-in-MultimodalonlineStoreEventDetection-in-EventBased-SocialNetworkSupport-A-
DistributionBased-AnomalyDetection-Methodfor-CloudendNetworkAn-EventDriven-
PrivacyPreserving-CollaborativeDetectionApproach-for-RankEvaluateAn-eventBased-Method-
for-MetamorphicEventSequenceMismatchingEventDetection-in-EventBased-
SocialNetworkSupport-DMossReactome-Event-Decting-Gluta-MatedWithoutDoubtspring-
using-NicMagnaMatin-linked-highImportantShschockInRabdomised-Control-TrialFor-
HyperbaricPrincipalsTwelveplus-One-A-time-Based-A-BayesianChainAimFor-
CircularGroupEvent-
EvidenceSaveMultilevelMultiAspectEventMiningInDarkWebClosedinHeuristic-
TheoryMultiAspect-TimeAwareEventDetectionModelwithAdaptiveSigmodEventRankonland-
UsingEventServicesNsseMoperation-basedSubEventDetectionDetectionMethodof-
IllegalSWOTAnalysisControlan-
EventHierarchicalDatabasesBasedZeroYlideRoutingMethodologyontheMetaOpenDataEventGpe
arDataverseGNumericalInvestigationof-
MemoryandDetectionforNonlocalDiffusionProblemSupport-A-Recommendation-Systemfor-

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



MA-w Technology-CompanybasedSocialMediaDatafor-SpecialtyProducts206-Event-Detection-A-DetectingOrientedAdaptiveAnomalyDetection-AlgorithmFormetalSinkDIgebouwzDPI-EventProcessingSupport-Fusionfor-DistancebasedEventPredictionToEnhanceEventComputationSupportDetectingAftersalesEventinCrossSellingDataSupportPolicyAnalysis-LTEValidationan-EventRelationDistributionBasedEventSeparationMethodBased-onDissectionProLate-TwInstatement-Detectionbib-DGMPDDE-Term-Detection-A-Dissertationon-SiliconEventservicesInduringvirtualDigital-PollinoseinSconnaEvent-Detection-a-Constrained-SATAWeiboEventRecognitionbased-onAuto-energy-A-PredictionModelwith-MultipleExpansionFunctionsforOpenInterestDEdiansocietyDynamicsDetectionUsingMultimodalControlMethods-DataLoveToLearn-A-GesturesupportedInteractiveSystemforLearningMathematicsinRemoteSENSORS-EventDetectionUsingMultiStageMultiTypeAttributes-A-Generic-EventStructureBasedEventDetectionExceptASupreventionof-FusionBasedSwarmRoboticsOversaturationanEventBasedEvaluationappl-N-A-Cascaded-LearningAlgorithmforReconizingHumanMovingDirectionand-EventDetectionwithoutSegmentDetectionofMovingEventsFrom-EventDataRecentOfAnEventDecisionRid-And-LearningIntelligentShelteredDitchEventDetectionMethodtemporal-CLAHEFor-EventDetectionFACE-Detectiona-FrameworkforMonitoringFishSchoolDemograph-EventTypeSelection--SemanticLocalization-A-DomainByDomen-TwoPhase-VerifiableSecureOutsourcingofFWAn-EffectiveEventDetectionModelinCENswith-RDTM-DetectionUsingEventCalibratedSupport-A-DistributedMultiCameraEventDetectionSchemaUsingNONEEventDataGuidedTHsmot-algorithmscores-Thek-MultiTypeEventDetectionAlgorithmbasedonDynamicCrowdsourcedGeo-SpecialEventData-Multimedia-EventDetection-ResponseEvaluationAcapACPTEventNetwork-detectionof-NotifiableDiseases-Using-NDMCNetworkDynamRunning-Based-Transitive-Closeness-FrequencySignalShift-DetectionFrameworkSupport-FusionofEvent-EventTypeDetectionUsingMultiClassJointEventAttention

Event-driven regulatory technology, or RegTech, refers to the use of an event-driven paradigm to support regulatory technology in the financial domain. In this context, events denote significant occurrences that are potentially relevant for compliance with laws or regulations, such as financial fraud, market manipulation and money laundering. Events are characterized by their value for law enforcement agencies, market watchdogs and regulatory bodies. Event-driven

RegTech is intended to facilitate the timely detection of violations of laws and regulations and is regarded as a response to the fast emergence of new-than-ever financial crimes.



Fig 2: Financial Crime

5. Research Summary

Intelligent detection mechanisms based on machine learning, statistical analysis, and graphs have been integrated into a complete architecture, and adapted to the detection of financial crime. The developed system is hosted on a private cloud of a financial institution, and

it generates and processes internal and external event streams. A new open-source package allows to define the entire detection process—from external entry points to monitoring dashboards—using only configuration files. The approach enables the emergence of complex and unexpected behavior from otherwise independent internal streams, and the implementation nature allows the overlay of additional detection methods as external modules. The concept of event processing has been further enriched by adding temporal reasoning, enabling the specification of logical rules that use time as a dimension of analysis for both simple and complex detection.

The combination of emerging event-driven paradigms with a well-known and widely used machine learning library dedicated to the financial field has allowed training and validation of bad behavior models by exploiting easily accessible public data, and testing against the latest consolidated data. Special attention was given to the selection of benchmark datasets and evaluation metrics, and the configuration of each detection method was executed in a systematic reproducible manner. Validated methods can therefore easily be reused by other parties interested in their implementation.

Equation B. Accuracy

Accuracy is also mentioned by the paper in evaluation.

Start with definition:

$$\text{Accuracy} = \frac{\text{all correct predictions}}{\text{all predictions}}$$

Correct predictions are $TP + TN$, and total predictions are $TP + TN + FP + FN$. Therefore:

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

5.1. Core Foundations of Event-Driven Regulatory Technology

The event-driven paradigm provides the foundational framework that addresses the primary principles of regulatory technology related to financial crime detection. Financial systems are inherently dynamic and continuously changing, characterized by an active

community of customers and employees whose transactions or customer support requests generate large amounts of data in the form of various events. The models and methods developed to detect, in real-time, financial crimes across all levels of financial institutions share some key aspects, regardless of how they function.

Event Processing: The combination of event streams, temporal reasoning, and complex-event processing underpin the rich event-driven paradigm. From the creation of an account to the closure of a suspicious account, the entire life cycle of a financial institution, including its employees, customers, accounts, transactions, and products, is event-based and event-centered. Events provide the highest abstraction level for describing the system states and system business functions. Stream reasoning through temporal intervals allows the derivation of new streams of evidence. Complex-event processing covers the detection of complex crimes based on the patterns of several events. Thus, the event-processing-layer capability is context-aware, dynamic, and suitable for detecting financial crimes.

Compliance Rules and Knowledge Base: The combination of rules-based and knowledge-base detection mechanisms, including the ability to combine all methods of intelligent detection in a layered structure, supports the corresponding nature of detection mechanisms. Financial systems are extremely complex, and crimes can occur at different levels with different degrees of sophistication. Regulatory authorities cannot impose excessive costs on banks. Hence, they use explicit, fixed, and quantifiable detection rules. Moreover, banks have the flexibility to formulate their own rules according to their specific procedures and customers. Banks also invest in detection models using advanced techniques, especially when there are information gaps in the regulatory authority. The logical combination of these detection methods allows for effective use of all detected evidence and fits the natural operating environment of any financial institution.

Table 2 — Detection methods


Detection method	Purpose in FinCompliance 2.0	Typical mathematical idea
Rule-based detection	Detect known suspicious patterns	Logical predicates / threshold rules
Temporal reasoning / CEP	Detect event sequences over time	Window functions, temporal predicates

Detection method	Purpose in FinCompliance 2.0	Typical mathematical idea
Anomaly detection	Detect rare deviations from normal behavior	Distance, density, or z-score style anomaly measures
Pattern recognition	Learn normal vs abnormal patterns	Classification functions
Network / graph-based compliance	Detect suspicious connectivity and subgraph structures	Graph degree, adjacency, path/subgraph analysis

6. Foundations of Event-Driven RegTech

The core principles of Event-Driven Regulatory Technology (EdRegTech) establish the foundation for the design and implementation of suitable basic architectures and building blocks, and are presented here using the concept of levels. The vision of EdRegTech consists of monitoring trends and events in real time to detect potential compliance violations using appropriately prepared data streams and applying rules, models, and algorithms from within a compliance knowledge base. The key overall idea revolves around the principle of data-driven compliance detection, which entails automatically discovering and monitoring data access and processing patterns, and detecting tranches of data or data processing patterns with non-compliance risk with respect to applicable regulations, or of suspicious data processing activities indicative of financial crimes. To facilitate the systematic development of building blocks and components fulfilling this vision, a two-level event processing model has been established.

At the lower level, compliance data streams are being prepared for compliance reasoning purposes, including temporal reasoning support for both compliance detection and event correlation. The temporal reasoning support enables the rationale of detecting potential compliance violations outside of the financial sector, while the prepared streams enable the application of complex-event-processing – and, in particular, event-detection – techniques. All these functions are key for enabling the intelligent detection and discovery of compliance violations that are hidden in the data, and for applying complex-event-processing techniques to directly process the compliance-detection reasoners. At the second level, the trends and events in the compliance data streams are spotted with the help of complex-event-processing technology and the applicability of the intelligent-detection paradigm is illustrated.



6.1. Core Principles of Event-Driven Regulatory Technology

Event-driven RegTech serves as a conceptual level modeling framework for regulatory technology solutions, encapsulating three key components: events, compliance rules, and an event processing environment. Events, including standard events predefined by regulators and system-generated events, function as triggers for executing compliance processes within the event processing environment. The combination of various event sources, processing paradigms, and algorithms creates a flexible and extensible infrastructure to support the intelligent detection of financial crimes.

FinCompliance 2.0, an implementation of event-driven RegTech for financial crime detection in banking, embodies these principles with a layered architecture comprising an event processing layer, a rules and knowledge base, intelligent detection mechanisms, and an evaluation layer. The event processing layer accepts various financial crime detection tasks and connects to event sources as an event consumer. Standard events for banking and money laundering, supported by Basel 3 and FATF recommendations, serve as a foundation, enabling regulatory authorities to bespoke a domain-specific compliance rules dataset. The refined event streams become augmentation sources for anomaly detection, complex event processing,

network-based detection, and graph-based compliance mechanisms, thereby ensuring detection completeness and accuracy.

7. Architecture of FinCompliance 2.0

FinCompliance 2.0 is designed as an event-driven system comprising an event-processing layer, a repository of compliance rules and knowledge, and intelligent techniques. The event-processing layer handles the ingestion and preliminary processing of events, such as filtering, enrichment, transformation, and probabilistic temporal reasoning through event streams, complex-event processing, and temporal-support-vector machines. The core processing of FinCompliance 2.0 resides in the compliance rules and knowledge base, which map different aspects of Financial Regulatory Compliance Records (FRCR) as defined by the European Union Wire Transfer Regulation (WTR). Finally, intelligent detection techniques based on anomaly detection and pattern recognition, network and graph-based compliance, and data-governance models are integrated into FinCompliance 2.0, enabling intelligent detection of financial crime while reducing false positives and covering additional detection scenarios.

Three data layers form FinCompliance 2.0. The first contains the events and data required for financial-crime compliance, such as transactions, customers, related actors, and security-violations records. The second layer consists of the compliance rules defined in the knowledge base. A third layer implements the intelligent-detection mechanisms for financial crimes such as money laundering, human trafficking, and terrorism financing.



Fig 3: RegTech's Role in Financial Compliance

7.1. Event Processing Layer

The event processing layer is at the heart of event-driven regulatory technology (RegTech). This functional component carries out Event Processing as a Service for RegTech-enabled applications (software that support compliance with regulatory requirements). Event Processing as a Service comprises components of data acquisition, including source message filtering, transformation, pre-processing, and enrichment, as well as support components. Event Processing as a Service processes messages and events that flow through the event-based architecture by exploiting data provenance, data lineage, and Complex Event Processing (CEP) for detecting financial crime. It uses an adaptation of the Event Processing Systems model and integrates recent advances in CEP to be fit in the specific context of financial compliance.

Message sources—inputs to Event Processing as a Service—that supply monitored data to applications and rule bases can be found in the data acquisition subsystem. Besides the data acquisition subsystem, support components (non-core workloads) can provide internal functions that relieve burden from the central Event Processing as a Service. Such components may cover data transfer layer activities (broadly managing) that are not performed by source filters and non-core persistence. Support not specifically required for the processing layer may also be provided by external applications connected to the communication layer.

Equation C. ROC curve and AUC

The refers to ROC curves and AUC scores in the results/validation discussion.

1. True Positive Rate

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

2. False Positive Rate

$$FPR = \frac{FP}{FP + TN}$$

3. ROC definition

An ROC curve plots TPR against FPR as the detection threshold changes.

So mathematically:

$$ROC = \{(FPR(\tau), TPR(\tau)) : \tau \text{ varies over thresholds}\}$$

4. AUC derivation

Area under the ROC curve is the integral of TPR with respect to FPR :

$$AUC = \int_0^1 TPR(FPR) d(FPR)$$

For discrete threshold points (FPR_i, TPR_i) , trapezoidal approximation is used:

$$AUC \approx \sum_{i=1}^{n-1} \frac{TPR_i + TPR_{i+1}}{2} (FPR_{i+1} - FPR_i)$$

7.2. Compliance Rules and Knowledge Base

The social science research domain has constructed a comprehensive knowledge base by compiling and formalising rules and relationships among concepts pertinent to human beings and society, enabling symbolic reasoning under the auspices of the Semantic Web paradigm. Similarly, the compliance rules and knowledge base in FinCompliance 2.0 consolidate governances related to financial service companies, governmental institutions, market players, and business partnerships into a shared knowledge structure, which is indispensable for effective temporal reasoning and event pattern discovery. More specifically, businesses in the financial domain are under regulatory supervision because they hold and manage the financial and personal property of citizens. Data-collecting companies such as banks and insurance companies collect and analyse large volumes of customer data over extended time periods to enhance investments and marketing strategies. The emerging trends and events detected in those databases can be shared with other market players in order to leverage business partnerships and optimise corporate performance. Nevertheless, the same features that motivate companies to collect such extensive data are exploited by financial criminals to perpetrate their malicious endeavours. Therefore, it is of utmost importance to detect emerging trends in data-represented behaviour over time so that financial crimes can be prevented.

FinCompliance 2.0 integrates established knowledge and temporal patterns extracted from historical data into the Compliance Rules and Knowledge Base. The temporal, pattern-learning engine of Complex Event Processing (CEP), combined with the Compliance Rule and Knowledge Base constructed at development time, supports the perception of recent trends in business data that deviate from normal activity, supports fraud-detection systems, and enables financial security management. Bank customers can be monitored to prevent identity fraud based on known behaviour by generating alarms whenever access is attempted from a different location or country without prior notification. The detection of anomalous transactions such as two large withdrawals in sequence can also trigger an alert. Suspicious money transfers to or from countries suspected of money laundering can be flagged. In addition, modelling a travel history extends these CPM-assisted capabilities to abnormal credit-card usage while travelling abroad.

8. Intelligent Detection Mechanisms

Intelligent detection mechanisms are crucial for automatically identifying complex financial crime schemes, especially when considering the multitude of different crime types, actors involved, and their interactions across multiple domains. Effective detection relies on implementing detection engines capable of several forms of exploratory analytics or machine learning: anomaly detection, pattern recognition, network-based compliance, and business rule validation in general.

Heavy reliance on standard business rules or empirical thresholds is ineffective or problematic due to the high volume of transactions, high dimensionality of the data, high variance in behavior, and class imbalance between “bad” and “good” classes**[6]**. Unsupervised anomaly detection techniques for alert generation can identify potentially suspicious transactions, subsequent graph analyses can determine which alerts originate from the same source or target, and finally, supervised models can rank the alert groups. The alerts are grouped, ranked, and routed based on a combination of the amount involved, network changes, and the risk characteristics of the involved entities.

8.1. Anomaly Detection and Pattern Recognition

A data-driven, machine-learning-based detection methodology supervised by domain experts provides further detection capabilities not achievable with strictly rule-based systems. Anomaly detection aims to detect a small number of rare observations that raise suspicions by differing significantly from the majority of data. It can be viewed in both a univariate and multivariate context. For instance, in the case of know-your-customer data, if the majority of customers are commonly derived from Europe and one customer is identified as American, rule-based systems will not flag this as suspicious. However, this American customer can very well be flagged as suspicious by the anomaly detection system because they differ from the majority of customers. Yet this should not be flagged as suspicious—after all, companies having American owners are not suspicious per se. This can also be explained by the fact that the majority of customers have given their residence outside of the United States, and the company operates in a sector dominated by Eastern Europe.

Pattern recognition is a method whereby a selection of normal behavior of the entities involved (for example, buyers, sellers, locations) is learned using traditional machine-learning algorithms, making it possible to then classify other operations as normal or abnormal since they

fall out of the learned model. However, these operations being classified as abnormal do not have to be labeled as suspicious; they are simply not in line with what has been learned. From a RegTech perspective, these two techniques could be trained by domain or data experts to allow them to be deployed using normal and semi-supervised learning.




Fig 4: 2025 Global Regulatory & Financial Crime Trends

8.2. Network and Graph-Based Compliance

The detection of criminal financial activity often benefits by explicitly testing for violations of business-related rules and compliance patterns. Detecting such irregularities using event-routing techniques has led to the identification of spokes in corporate networks, expression of model management and control in bounding-key policies, and steering of capital movements using temporal transaction-processing languages. Nevertheless, the most powerful form of such analysis utilizes some aspect of subgraph matching in the underlying connection graph.

Steering compliance record compliance with numerous business rules without resorting to the full complexity of graph queries is made possible by expressing certain kinds of guard predicates directly in key-based conditions. Such guard predicates usually take the form of existence tests: Is there a licensed operational subsidy whose business nature is detected for the conversant country, its country of activity not coinciding with that of the license?

9. Event-Driven Processing Paradigms for RegTech

Event-driven architectural paradigms, such as event streams and complex event processing, offer substantial benefits to the field of regulatory technology. Event streams can act as an abstraction over sub-components of a RegTech solution—for instance, the transaction ledger of a financial institution. They can also provide a foundation for additional regulatory capabilities, such as auditing. Complex Event Processing (CEP) in particular provides a foundation for regulation-aware supervisory monitoring and financial crime detection and takes place at the logical level closest to the border of the RegTech solution. Its support for temporal reasoning can significantly enhance financial crime detection affecting a particularly broad scope of crime, that of money laundering. Brightly lit in its sweet spot, CEP can face challenges such as scalability, accuracy, and recall with diminishing overhead.

An event-driven architecture is a natural choice for implementation of any intelligent financial detection system, considering the high volume and velocity of transactional data. At first glance, Regulation Technology (RegTech) appears a niche concept, yet any mature financial institution should have a module or toolkit focusing on financial legislation compliance—either in-house or outsourced. Building a solution from the ground up offers benefits of knowledge integration and offering priced services for other firms. At deeper inspection, event-driven architectures are ideally suited for Regulatory Technology, exploiting intuitive sub-components edges without requiring a guard rail around the sweet spot.

9.1. Event Streams and Temporal Reasoning

Complex event processing (CEP) utilizes event streams—or continuous data flows that occur over time and can be processed incrementally—as a general model. The model is well suited to financial crime detection due to the challenge of monitoring activities at scale. Financial operations occur continuously and must be evaluated quickly to detect and make sense of anomalies. Timeliness and order are intrinsic characteristics of financial operations; therefore, a

CEP approach serves as a natural means of organizing and processing operations in compliance detection pipelines.

Building on previous work, an event-stream-based processing mechanism is proposed for financial crime detection. Each data stream is represented as a sequence of data items ordered by timestamp, which can be continuously processed in a windowing manner. Stream alignment provides a basis for querying events with temporal semantics; the detected temporal patterns of events act as the foundation for assurance reasoning. In practice, compliance-related queries focused on temporal aspects can be expressed as CEP queries and evaluated over event streams.

9.2. Complex Event Processing in Financial Crime Detection

Financial crime detection inherently requires progressively more sophisticated data analysis and detection techniques as criminals become more adept and organized and schemas become more complex. In addition to Boolean algebraic rules based on observations or knowledge, complex event processing (CEP) has recently become a powerful tool to detect criminal events. In CEP systems, events are derived from sensor data and sent to event processing engines where they can undergo "event pattern matching" in near real time.

Detection strategies based on CEP respond to enabling technologies that offer smarter, faster and cheaper handling of massive amounts of event data: (1) the vast amounts of event data generated by sensor technology, (2) the need to process that data rapidly, (3) the growing interest in pattern recognition, and (4) financial institutions' desire to look beyond simple anomaly detection and realize intelligence-driven solutions. In criminal detection, pattern queries for known and believed behaviours by experts represent knowledge accumulated during investigations or monitoring—a huge value that can possibly be further developed using machine learning. The knowledge model should support both detection mechanisms (anomaly detection and pattern matching) and retain the knowledge accumulated over the divisions' years of operations.

Table 3 — Metrics explicitly or implicitly

Metric	Formula	Meaning
Precision	$\frac{TP}{TP + FP}$	Out of predicted positives, how many were truly suspicious



Metric	Formula	Meaning
Recall	$\frac{TP}{TP + FN}$	Out of all true suspicious cases, how many were detected
F1-score	$\frac{2PR}{P + R}$	Balance of precision and recall
TPR	$\frac{TP}{TP + FN}$	Same as recall
FPR	$\frac{FP}{FP + TN}$	Fraction of normal cases wrongly flagged
AUC	$\int_0^1 TPR(FPR) d(FPR)$	Overall ROC performance
Accuracy	$\frac{TP + TN}{TP + TN + FP + FN}$	Total correctness

10. Data Governance, Privacy, and Compliance

The management of sensitive information in accordance with legal and organizational requirements for data governance and compliance with privacy regulations must be considered in the development of a framework for intelligent financial crime detection. Therefore, the proposed solutions should support data lineage, provenance, and privacy-preserving data sharing and publishing mechanisms.

Data lineage enables tracking the life cycle of individual data elements across their transformation, derivation, and movement within a system, while supporting the identification of responsible parties for data quality and compliance. Compliance with the EU General Data Protection Regulation demands provision for the data subject's right to erasure and other privacy obligations, as well as assessment of the potential privacy impact of administrative activities. Data provenance refers to the automatic documentation of the history of data products for trust and quality assessment and must record the source, process, and transformation of data products in sufficient detail. Furthermore, privacy-preserving data sharing is critical for the prevention of financial crimes, with techniques including anonymization, k-anonymity, k-degree anonymity, and differential privacy applicable to the publication of anonymized datasets and location data.

Equation D. Event-stream representation and temporal windowing

The emphasizes **event streams**, **timestamps**, **temporal reasoning**, and **CEP**.

1. Event stream

Represent the stream as a time-ordered sequence:

$$S = \{e_1, e_2, e_3, \dots, e_n\}$$

Each event has attributes and a timestamp:

$$e_i = (x_i, t_i)$$

where x_i is the event payload and t_i is the event time.

2. Time ordering

For event streams:

$$t_1 \leq t_2 \leq t_3 \leq \dots \leq t_n$$

3. Sliding time window

Suppose we inspect the last W time units at time t .

Then the active window is:

$$\mathcal{W}(t) = \{e_i \in S \mid t - W \leq t_i \leq t\}$$

4. Count of suspicious events in the window

Let indicator $I(e_i)$ be:

$$I(e_i) = \begin{cases} 1, & \text{if } e_i \text{ satisfies a suspicious predicate} \\ 0, & \text{otherwise} \end{cases}$$

Then total suspicious events in the window are:

$$N_{\text{susp}}(t) = \sum_{e_i \in \mathcal{W}(t)} I(e_i)$$

5. Rule trigger

If the threshold is θ , then a temporal rule can be written as:

$$N_{\text{susp}}(t) \geq \theta \Rightarrow \text{Alert}$$

10.1. Data Lineage and Provenance

Deterministic and automated Compliance-as-a-Service approaches yield difficulties in compliance validation and risk management due to lack of transparency and guaranteed reliability. Analogous to the role of data lineage and provenance in Business Intelligence and Big Data domains. Tracking the data provenance provides a temporally and contextually grounded enterprise-wide overview of the data sources used in these systems.

Transfer learning techniques utilize historical data as a surrogate for real-time commercial and regulatory data, bypassing complexities related to event streams and temporal reasoning. However, increased correlations pose difficulty. Metadata capturing key features of each dataset – including sampling period, sampling rate, target distribution, and properties of the transfer learning surrogate data – mitigate risk, enhance interpretability, and test support requirements of potential transfer learning solutions.

10.2. Privacy-Preserving Techniques

Privacy threats are inherent in the financial sector due to the collection and processing of sensitive data like Personally Identifiable Information (PII). Existing privacy-preserving methods fall into three categories: anonymization, cryptographic techniques, and privacy-preserving data mining. Data anonymization protects sensitive information by replacing original data with fictitious or non-meaningful data by techniques like generalization, k-anonymity, l-diversity, t-closeness, blind data out-sourcing and the use of data masking algorithms, etc. Data encryption produces secure data but can hinder the performance of detection systems. Additionally, these methods do not preserve semantic information, which is often crucial for effective financial-

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



crime detection. Privacy-preserving data mining techniques mainly concern using privacy-preserving data mining protocols to protect sensitive data during the data mining process.

Differential privacy has been proposed as a solution to these problems in the context of data publishing. The approach adds suitable noise to the sensitive data and offers strong guarantees regarding potential disclosure risks. The concept of differential privacy may be applied to different levels of the system architecture, including the detection and learning phases. For instance, a privacy-preserving K-means clustering algorithm allows sensitive data to remain private while still enabling the use of K-means clustering. The privacy-preserving data classification technique based on multiple parties combines L-diversity, Non-Interactive Zero Knowledge Proof, and Homomorphic Encryption in decision tree classification. Such techniques allow different parties to build a decision tree classifier without sharing their private data with each other.

Data owners control access to encrypted data by means of special authority, which enables fine-grained access based on user identity. Sensitive data is encrypted, but the intention is to enable public auditing of the data integrity. To address possible data leakage problems, data owner confidentially outsources the sensitive data and provides fine-grained access to authorized users. Another technique provides both data integrity and confidential authorization using an asymmetric weak key system. Users can monitor the integrity of the data and get only the required parts of the data without expending extra costs.





Fig 5: Major Perks of RegTech Implementation

11. Evaluation and Validation

A comprehensive assessment of the proposed event-driven RegTech framework addresses data governance, privacy, and compliance implications. Key aspects include establishing data lineage and provenance and applying privacy-preserving techniques together with a dedicated lens on legal and regulatory facets. The second part of the validation leverages publicly available datasets and metrics to benchmark financial crime detection and compliance capabilities against state-of-the-art tools.

To fulfil the frameworks promise of supporting consequential decision-making, detection effectiveness is evaluated using a variety of specific objectives—detection accuracy and F1 score for supervised models, intrinsic detectability for unsupervised models, control-scoring capabilities in network analytics, and compliance-preserving mechanism matching in graph-

based methods. Validation shows performance of supervised classifiers suited to compliance supervision, with partial support for anomaly detection from hold-out F1 score and ROC-AUC surface associations. Unsupervised techniques offer intrinsic latent-group discovery, and the knowledge-based event model controls spatiotemporal crime network behaviours while acknowledging resource limitations in high-dimensional data analysis. Context-driven intertwining of privacy, data governance, and detection further advances detection paradigms.



11.1. Benchmark Datasets and Metrics

Event-driven technology encompasses a variety of techniques that are well-suited for continuous, lightweight, and horizontal processing of large streams of events, as shown in the myriad of implementations in application domains such as financial crime detection and prevention, finance market surveillance, and network traffic monitoring. Completeness and

correctness are clearly important requirements to avoid the filtering or omission of meaningful data. The evaluation of detection and processing capabilities can also take into consideration the concordance by comparison with other complementary detection components. An appropriate benchmark dataset, together with detection metrics, is essential in the assessment process.

The publicly-available, benchmark event stream dataset for credit card fraud detection contains 30 months (approximately 1.5 billion events) of real-world transactions made by European credit cards. An extension of the dataset with split timestamps, labels, and remapped category codes is suitable for binary supervised and ordinal supervised anomaly detection. The extension of the dataset also includes additional temporal control information for temporal reasoning support—an essential requirement of credit card fraud detection systems, which are designed to filter transactions as real-time events in a 24-hour cycle. The detection system achieves high fidelity and precision measurement values for this dataset. Performance evaluation of two autoencoder-based fraud detection solutions and the respective schema and detected categories yield satisfactory fidelity, true positive, and false positive rates. The detection performance of the solutions delve into data residency and privacy concerns for image-processing networks and their corresponding effects.

The detection metric concordance investigation considers the performance of three external network-consistent constraint-based solutions that analyse and connect temporal-spatial trajectory networks. Evaluation against the Cyber Threat Intelligence (CTI) database populates two additional other-operator labels, with detected C2 clusters contributing a further label. A proprietary network-attack benchmark dataset covers attacks such as DNS (Domain Name System) flooding, grey-hole, map-reduce, and forwarding-loop, and non-attack scenarios at the IP layer.

Equation E. Z-score anomaly detection

The anomaly detection as a mechanism for detecting rare deviations.

Suppose one feature is transaction amount x .

1. Mean

For n observations $x_1, \dots, x_n$,



$$\mu = \frac{1}{n} \sum_{i=1}^n x_i$$

2. Standard deviation

$$\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2}$$

3. Z-score

For a new transaction x ,

$$z = \frac{x - \mu}{\sigma}$$

4. Absolute anomaly score

Often suspiciousness is based on distance from the mean:

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

5. Thresholding

If the threshold is k ,

$$|z| > k \Rightarrow \text{anomaly}$$

11.2. Validation of Detection Effectiveness

The effectiveness of detection mechanisms incorporated into FinCompliance 2.0 is evaluated using well-established benchmark datasets for AML and fraud detection. The models have been trained or developed utilizing the benchmark datasets, and performance metrics from

the relevant literature, ensuring that detection effectiveness, accuracy, and efficiency are rigorously validated. The monitoring of financial transactions is not restricted to the fields or attributes defined by the data model used to build the underlying rule set. Temporal analysis for FinCompliance 2.0 can also be performed using the complex-event- and rule-based capacity of event-processing systems built around CEP engines. Event data linkage with open-source projects such as the Open Information Security Foundation's Open Source Malware Information Sharing Platform, taxonomies created by MITRE, the Cyber Threat Intelligence Reference Architecture, and the European Union Agency for Cybersecurity provides an enhanced information base that can also be used for broad-temporal-security analysis. In the area of NLP, new developments are introduced, for example, for textual document similarity.

Event-processing technologies create a facility for implementing and processing event data streams, utilizing data with defined and logical temporal evolution as explored in the Open Data and Windows of Opportunity projects. These techniques, supported by Big Data processing, allow the combination of event streams of different natures and sources, feeding war-gaming and simulation systems that exploit the data in real time as if a physical security incident were happening.

12. Results

The effectiveness of the proposed solution, prototype architecture, and intelligent detection capabilities was evaluated and validated. Established benchmark datasets were employed to verify the accuracy and performance of the key supervisory detection mechanisms. The results demonstrate the suitability of the approaches for real-time compliance monitoring through event streams, temporal reasoning, and complex event processing paradigms.

To assess the detection quality of the systems and their ability to capture compliance requirements, reference real-world datasets originating from the credit card and money laundering domains were used. The performance of the key detection functions was evaluated by evaluating the ROC curves and AUC scores.




Fig 6: Regtech Market for Financial Crime Compliance

13. Conclusion

Research into the event-driven RegTech paradigm led to the development of FinCompliance 2.0, represents a reiteration and evolution of FinCompliance 1.0, focuses on intelligent detection for complex financial crimes, and explores the detection of illicit financial activities in a holistic manner. Leveraging event-driven computing methodologies helps to unify the detection of sophisticated financial crimes and real-world cyber-criminal activities that are spatially and temporally related by integrating detection capabilities into a central infrastructure component. Moreover, appropriate detection techniques for a specific set of events can be incorporated in the central infrastructure in an as-needed manner. Accordingly, potential detection capabilities span temporal reasoning and event stream processing, combined with supplementary data, to achieve explicit detection objectives. The architecture was augmented with methods for the temporal reasoning layer and the complex event processing layer. The research further enhances the event-driven RegTech paradigm by identifying detection techniques that support discovery. Internal provenance and lineage mechanisms augment data attribute and value auditing.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



As a real-time detection framework, FinCompliance 2.0 adopts an event-processing layer with full event schema support, a compliance knowledge base to store compliance rules, and a detection engine to execute the compliance rules to regulate abnormal events or behaviours. The static event schema constitutes a rich source for anomaly-detection-based and pattern-recognition-based detection mechanisms. The event-processing layer also enables flexible detection methods. If additional external data are available, information from the data sources may also be used for detection. The event-processing layer monitors event streams in real time to provide a triggering mechanism for detection without cumbersome resource-consumption on unprocessed data.

References

1. Alarab, I., & Prakoonwit, S. (2023). Graph-based LSTM for anti-money laundering: Experimenting temporal graph convolutional network with Bitcoin data. *Neural Processing Letters*, 55, 689–707.
2. Astrova, I. (2023). Anti-money laundering powered by graph machine learning: “Show me your friends and I will tell you who you are.” *Intelligent Decision Technologies*, 17(1), 243–261.
3. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. *International Journal Of Finance*, 36(6), 682-706.
4. Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 8, Article 151.
5. Bhowmik, A., Sannigrahi, M., Chowdhury, D., Dwivedi, A. D., & Mukkamala, R. R. (2022). DBNex: Deep belief network and explainable AI based financial fraud detection. In 2022 IEEE International Conference on Big Data (Big Data). IEEE.
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. (2022). Implementing Databricks Unity Catalog For Centralized Data Governance In Multi-Business-Unitenterprises. *Journal of International Crisis and Risk Communication Research*, 101-122.



8. Cheng, D., Ye, Y., Xiang, S., Ma, Z., Zhang, Y., & Jiang, C. (2023). Anti-money laundering by group-aware deep graph learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(12), 12444–12457.
9. Eddin, A. N., Bono, J., Aparício, D., Polido, D., Ascensão, J. T., Bizarro, P., & Ribeiro, P. (2022). Anti-money laundering alert optimization using machine learning with graphs. *arXiv*.
10. Gubran Al-Hashedi, K., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, 100402.
11. Mattaparthi, R. (2023). Deep Learning-Driven Combustion Anomaly Detection in Diesel Powertrains: A Multi-Sensor Fusion Approach for Real-Time ECM Adaptation. *International Journal of Intelligent Systems and Applications in Engineering*, 11, 1084.
12. 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.
13. Jensen, R. I. T., & Iosifidis, A. (2023). Qualifying and raising anti-money laundering alarms with deep learning. *Expert Systems with Applications*, 214, 119037.
14. Johannessen, F., & Jullum, M. (2023). Finding money launderers using heterogeneous graph neural networks. *arXiv*.
15. Kolla, S. H., & Loganathan, R. (2023). Cloud-Native Deep Learning Architectures For Secure Generative AI Deployment In Enterprise Workflow Platforms. *Journal of International Crisis and Risk Communication Research*, 603-618.
16. Kim, D.-H., Lee, S.-G., & Jung, S. K. (2021). Design and implementation of a deep learning model for fraud detection system (FDS). *Journal of Convergence Security*, 21(5), 69–78.
17. Kumar, A., Prusti, D., Das, D., & Rath, S. K. (2021). Fraud detection in anti-money laundering system using machine learning techniques. In *Proceedings of the International Conference on Paradigms of Computing, Communication and Data Sciences* (pp. 687–699). Springer.
18. Lokanan, M. E. (2022). Predicting money laundering using machine learning and artificial neural networks algorithms in banks. *Journal of Applied Security Research*, 19(1), 20–44.
19. Reddy, V. A. R. (2022). Designing Fault-Tolerant Data Ingestion Pipelines for High-Volume Healthcare Transactions. *Frontiers in Health Informatics*, 11, 861-889.
20. Lokanan, M. E. (2023). Predicting mobile money transaction fraud using machine learning algorithms. *Applied AI Letters*, 4(2), Article e85.
21. Rocha-Salazar, J.-d.-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, 114470.

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 26

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 22

ISSN: 3067-1140



22. Sannidhanam, A. H. (2021). Real-time claims processing using event-driven architectures. *International Journal of Technology, Management and Humanities*, 7(1), 36–50.
23. Syed, S. (2023). Shaping The Future Of Large-Scale Vehicle Manufacturing: Planet 2050 Initiatives And The Role Of Predictive Analytics. *Nanotechnology Perceptions*, 19(3), 103-116.
24. Stojanović, B., Božić, J., Hofer-Schmitz, K., Nahrgang, K., Weber, A., Badii, A., Sundaram, M., Jordan, E., & Runevic, J. (2021). Follow the trail: Machine learning for fraud detection in fintech applications. *Sensors*, 21(5), 1594.
25. Tuffveson Jensen, R. I., & Iosifidis, A. (2023). Qualifying and raising anti-money laundering alarms with deep learning. *Expert Systems with Applications*, 214, 119037.
26. Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing*, 149, 110984.
27. Mangalampalli, B. M. (2023). AI-Driven Anomaly Detection in Healthcare Claims Data: A Business Intelligence Perspective. *Journal of Rare Cardiovascular Diseases*.
28. Verma, T., & Misra, A. (2023). Financial fraud detection in financial institutions using two-layer-deep learning and self-improved honey badger algorithm. *Journal of International Finance and Economics*, 23(3), 30–54.
29. Wei, T., Zeng, B., Guo, W., Guo, Z., Tu, S., & Xu, L. (2023). A dynamic graph convolutional network for anti-money laundering. In *International Conference on Intelligent Computing* (pp. 493–502). Springer.
30. Wu, B., Lv, X., Alghamdi, A., Abosag, H., & Alrizq, M. (2023). Advancement of management information system for discovering fraud in MasterCard based intelligent supervised machine learning and deep learning during SARS-CoV-2. *Information Processing & Management*, 60(2), 103231.
31. Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. *International Journal of Scientific Research and Modern Technology*, 1(12), 227-237.
32. Wu, R., Ma, B., Jin, H., Zhao, W., Wang, W., & Zhang, T. (2023). GRANDE: A neural model over directed multigraphs with application to anti-money laundering. *arXiv*.
33. Zhou, X., & colleagues. (2023). Who are the money launderers? Money laundering detection on blockchain via mutual learning-based graph neural network. In *2023 International Joint Conference on Neural Networks (IJCNN)*. IEEE.