



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IX **Month of publication:** September 2026

DOI: <https://doi.org/10.22214/ijraset.2026.84957>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

AI-Driven Financial Fraud Detection: An Integrated Framework for Explainability, Privacy, and Resilience

Dr. S. Uma¹, R. Karunambigai²

Department of Computer Science and Engineering, Hindusthan College Of Engineering and Technology, Coimbatore, India

Abstract: *This paper presents a comprehensive framework for modern financial fraud detection that integrates machine learning, deep learning, graph-based learning, behavioral biometrics, explainable artificial intelligence, federated learning, adversarial defenses, and agentic AI. The study examines how ensemble models such as XGBoost, LightGBM, CatBoost, and stacking architectures support high-volume transaction analysis, while recurrent and graph neural networks capture temporal and relational patterns associated with sophisticated fraud schemes. Behavioral biometrics is considered for continuous authentication, whereas agentic AI is discussed for automated triage, investigation, suspicious activity reporting, and incident-response workflows. The paper also addresses asset tracing and recovery, victim notification, model explainability using SHAP and LIME, privacy-preserving federated learning, and adversarial attacks against fraud detection systems. Finally, regulatory considerations, including the EU AI Act, and future directions such as neurosymbolic AI are examined. The framework emphasizes the need for accurate, transparent, privacy-aware, resilient, and accountable fraud detection systems capable of supporting continuous financial crime prevention and compliance.*

Keywords: *Artificial intelligence, financial fraud detection, machine learning, deep learning, graph neural networks, behavioral biometrics, explainable AI, federated learning, agentic AI, adversarial attacks*

I. INTRODUCTION

The worldwide financial system is presently undergoing a fundamental shift from responsive, regulation-driven supervision to anticipatory, autonomous protective frameworks. This evolution is necessitated by the staggering scale of financial crime, with global estimates suggesting that organizations lose approximately 5% of annual revenues to fraud, a figure equivalent to trillions of dollars.[1, 2] As digital transactions in the United States alone exceeded 2.7 billion in 2023, the sheer volume and velocity of data have rendered traditional manual reviews and static threshold-based systems largely inadequate.[1, 3] Modern fraud detection now relies on a sophisticated hierarchy of artificial intelligence (AI) and machine learning (ML) techniques designed to identify complex patterns, adapt to evolving threats in real time, and manage the entire lifecycle of an incident—from initial anomaly detection to the final submission of regulatory reports and the coordination of international asset recovery.[4, 5, 6]

II. METHODOLOGICAL FOUNDATIONS OF FRAUD IDENTIFICATION

The architectural backbone of contemporary fraud detection consists of a multi-layered approach that integrates classical machine learning with advanced deep learning (DL) models. The primary objective is to differentiate between legitimate consumer behaviour and fraudulent interventions with millisecond-level latency while maintaining high recall to minimize the financial and reputational damage of missed fraud.[4, 5, 7]

A. Gradient Boosting and Ensemble Architectures

In the current landscape, tree-based ensemble models remain a cornerstone for tabular transaction data due to their high predictive accuracy and relatively lower computational overhead compared to deep neural networks.[3, 8] Models such as XGBoost, LightGBM, and CatBoost are frequently deployed in "stacking ensembles," where multiple models serve as base learners and their outputs are fed into a meta-learner, typically a logistic regression or a multi-layer perceptron, to refine the final classification.[9] This approach is particularly effective at reducing variance and bias, which are inherent in the noisy environments of financial systems.

The performance of these models is often optimized using Bayesian hyper parameter tuning, which efficiently explores the parameter space to maximize metrics such as the Area Under the Precision-Recall Curve (PR-AUC) and the F1-score.[9, 10] In empirical evaluations using the IEEE-CIS Fraud Detection dataset, which contains over 590,000 real-world records, stacking ensembles have achieved accuracy and AUC-ROC scores as high as 99%, significantly outperforming traditional decision trees and linear models.[9, 11]

TABLE I
COMPARISON OF MACHINE LEARNING ARCHITECTURES FOR FINANCIAL FRAUD DETECTION

Model Architecture	Structural Approach	Optimization Focus	Transaction Domain Suitability
XGBoost	Gradient boosted decision trees with L1/L2 regularization	Minimizing classification error and computational cost	High-volume retail transactions [3, 9]
LightGBM	Leaf-wise tree growth with histogram-based algorithms	Training speed and memory efficiency in large-scale data	Real-time payment processing [3]
CatBoost	Ordered boosting to handle categorical features	Accuracy without extensive feature pre-processing	Insurance claims and loan applications [9, 12]
Stacking Ensemble	Hierarchical multi-model integration	Generalization across diverse fraud patterns	Enterprise-level cross-border monitoring [9]

While these models are highly effective for structured data, they often fail to capture the long-term temporal dependencies that characterize sophisticated money laundering or account takeover schemes. This limitation has driven the adoption of deep learning architectures that can model transactions as a continuous time-series.

B. Deep Learning and Sequential Analysis

Deep learning architectures excel in capturing non-linear dependencies and temporal patterns that escape tree-based models.[13, 14] Convolutional Neural Networks (CNNs) are often applied to transaction "images" or reshaped tabular data to identify spatial correlations, while Recurrent Neural Networks (RNNs) and their variants are used for sequential analysis.[1, 13]

Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRU) are particularly effective for modelling the "velocity" and "rhythm" of transactions. Fraud is rarely an isolated event; it often manifests as a sequence of actions—such as a small "ping" or authorization check followed by a rapid series of high-value transfers.[7, 15] LSTMs utilize memory cells and gating mechanisms to maintain long-term dependencies, allowing the system to recognize when a current transaction deviates from a user's established historical trajectory.[13, 16]

The mathematical representation of the LSTM update gate mechanism is critical for understanding its role in temporal feature retention. The system calculates an input gate it , a forget gate ft , and an output gate ot at each time step t to determine which information from the previous state $ht-1$ and current input xt should be preserved:

$$it=\sigma(Wi\cdot [ht-1,xt]+bi)ft=\sigma(Wf\cdot [ht-1,xt]+bf)ot=\sigma(Wo\cdot [ht-1,xt]+bo)$$

These gates ensure that the model selectively "remembers" relevant historical spending patterns while "forgetting" transient noise, such as one-off vacation spending.[16] Hybrid models, such as the Temporal-aware Heterogeneous Graph Oversampling and Attention Fusion Network (THG-OAFN), combine these sequential capabilities with Graph Neural Networks to process both time-series data and the complex topological relationships between financial entities.[16]

C. Relational Intelligence via Graph Neural Networks

A significant advancement in fraud detection involves the use of Graph Neural Networks (GNNs) to model the financial system as a heterogeneous graph where nodes represent entities, such as accounts, merchants, or IP addresses, and edges represent transactions or shared attributes.[2, 4, 16] Unlike traditional models that analyse transactions in isolation, GNNs capture "guilt-by-association," which is vital for identifying organized crime rings and money laundering networks.[2, 17]

GNNs operate by aggregating information from a node's neighbours, allowing the model to learn a representation that includes the context of its connections. If an account is linked to a known fraudulent IP address or interacts with a confirmed "money mule," the GNN propagates this risk signal through the network.[2, 17] Techniques such as Adaptive Sampling and Aggregation-Based Graph Neural Networks (ASA-GNN) address the limitations of standard GNNs by focusing on the most relevant neighbouring nodes and employing entropy-based aggregation to combat "oversmoothing," where legitimate and fraudulent nodes become indistinguishable in the embedding space.[2]

TABLE II
GRAPH-BASED FEATURES FOR FINANCIAL FRAUD DETECTION

Graph Feature Category	Specific Metric	Fraud Detection Utility
Node Attributes	Account age, historical risk score, KYC status	Assessing baseline entity reliability [17]
Edge Properties	Transaction amount, timestamp, currency	Identifying unusual flow dynamics [16]
Topological Metrics	Degree centrality, PageRank, community membership	Spotting hub-and-spoke money laundering structures [17]
Temporal Signatures	Transaction velocity, arrival time distribution	Detecting automated bot or script behaviour [16, 17]

Furthermore, the GRAN (Graph-based Relational Anomaly Network) framework allows for the extraction of these topological features without the direct computational cost involved in standard GNNs, making it more practical for real-time deployment in high-intensity financial environments.[17] These features are then integrated into weighted ensembles to improve the detection of complex relational patterns.

D. Behavioural Biometrics and Neurobiological Baselines

The shift from static credentials, such as passwords and PINs, to behavioural biometrics represents a move toward continuous authentication. AI systems now verify identity based on *how* a person interacts with their device, creating a neurobiological unique profile that is significantly more difficult for fraudsters or bots to replicate.[7, 15, 18]

Behavioural fraud detection examines multiple layers of user activity to establish a baseline of "normal" behaviour. Deviations from this baseline trigger alerts even if the transaction itself appears legitimate based on static rules.[15, 19] These systems can achieve an accuracy rate of 98.7% against synthetic identity fraud and deep fake attacks.[7]

- 1) **Keystroke Dynamics:** The system measures typing speed, the rhythm of key presses, and the pressure applied to keys. Bots and fraudulent scripts typically exhibit robotic, perfectly consistent timing that deviates from human variability. [7, 18, 19]
- 2) **Mouse and Navigation Patterns:** This involves tracking mouse trajectories, click speeds, and scrolling behaviour. Human mouse movements are rarely perfectly linear; they contain micro-hesitations and curves that reveal human intent.[7, 18]
- 3) **Touchscreen Interactions:** On mobile devices, AI analyses swipe gestures, the angle of the device (orientation), and the surface area of the finger touch. These patterns are unique to an individual's physical handling of the phone.[18, 19]
- 4) **Facial Micro-Expressions:** Advanced systems can detect involuntary movements such as blink rates and mouth tension. These reveal neurobiological authenticity that current deepfake technology often lacks. [7]

By analysing these cues in real-time, banks can detect "mule accounts" where the legitimate owner is being manipulated or where the account has been entirely taken over by an automated script. These systems move from asking "Is this the correct password?" to "Is this the behaviour I expect from this user, in this context, at this time?" [7, 20]

The "bottleneck" in modern fraud operations is often not the detection of suspicious activity, but the subsequent investigation and reporting process. [21, 22] Traditionally, writing a single Suspicious Activity Report (SAR) can take an experienced analyst anywhere from 25 to 35 minutes. [22] AI is now being deployed to manage this entire lifecycle, transforming from a simple detection tool into a "continuous compliance auditor". [5]

III. AGENTIC AI IN INVESTIGATION AND TRIAGE

"Agentic AI" refers to systems that do not just flag anomalies but autonomously execute multi-step workflows. These agents handle the Level 1 (L1) triage that historically consumes 60% to 80% of an analyst's time.[21] When a transaction is flagged, the AI agent pulls transaction histories, checks device signals against shared fraud intelligence networks, gathers evidence from multiple data sources, and assembles a case narrative.[21, 23]

TABLE III
HUMAN-AI AGENT ROLES IN THE FINANCIAL FRAUD DETECTION WORKFLOW

Workflow Stage	Human Analyst Role	AI Agent Role
Detection	Policy setting and rule oversight	Real-time transaction evaluation and risk scoring [5, 21]
Triage	Reviewing high-risk edge cases	Pulling evidence, checking watch lists, and L1 disposition [21]
Investigation	Complex entity linking and judgment calls	Graph analysis, OSINT research, and narrative drafting [21, 23]
Reporting	Final approval and regulatory signing	Auto populating SAR fields and generating narratives [22, 24]

E. Post-Detection Interventions and Incident Response

Once an AI system identifies a high-confidence fraud signal, it triggers a cascade of actions designed to mitigate immediate financial loss and initiate the recovery process. These actions range from automated technical blocks to multidisciplinary legal and psychological support for victims.[5, 25, 28]

F. Real-Time Mitigation and Step-Up Authentication

The first line of defence is often a "silent" action that protects the user without disrupting their experience unless necessary. If an AI identifies a transaction as potentially fraudulent, it might assign a high risk score that triggers a tiered response:

- 1) Transaction Blocking: For high-confidence fraud, the payment is halted immediately before it is authorized by the network.[5, 19, 25]
- 2) Account Freezing: In cases of suspected account takeover, the system freezes the entire online banking profile to prevent further charges.[18, 25, 29]
- 3) Multi-Factor Authentication (MFA) Step-Up: If the risk is moderate, the system may prompt the user for additional verification, such as a biometric scan or a secure access code, to confirm their identity.[5, 18, 29]

These automated actions significantly reduce the window of opportunity for fraudsters to move funds out of an account. In finance operations, these AI agents can also initiate workflows that request supporting documentation from the user or merchant, accelerating the resolution of policy exceptions. [5, 21]

IV. EXPLAINABLE AI (XAI) AND THE TRUST EQUATION

The transition to complex "black-box" models, such as deep neural networks, has raised concerns regarding trust, accountability, and regulatory compliance.[12, 36, 37] Explainable AI (XAI) is now a prerequisite in high-stakes financial domains where decisions must be justified to both regulators and affected individuals.[8, 38, 39] XAI techniques are categorized into "ante-hoc" models, which are inherently interpretable (like decision trees), and "post-hoc" models, which are applied to deep learning systems after a prediction is made.[36]

- 1) SHAP (SHapley Additive exPlanations): Based on cooperative game theory, SHAP determines the contribution of each input feature to the final prediction. In credit card fraud, it might identify that the "transaction amount" and "merchant category" were the primary drivers for a high risk score.[8, 9, 40] The "SHAP Stability Index" is used to measure the consistency of these explanations across different datasets, ensuring that the AI isn't relying on transient noise.[10]
- 2) LIME (Local Interpretable Model-Agnostic Explanations): LIME provides transparency by approximating a complex model's behavior in the vicinity of a specific transaction. This allows an investigator to see exactly why a single alert was triggered, making the output of a deep learning model auditable and understandable to a non-technical user.[9, 37, 40]
- 3) Counterfactual Explanations: These provide actionable insights by showing the minimal change needed for a different decision (e.g., "If the transaction had occurred from a known IP address, it would not have been flagged").[36, 37]

These tools are essential for aligning AI performance with organizational accountability mechanisms and complying with directives like the GDPR's "Right to Explanation".[14, 37, 40]

V. FEDERATED LEARNING: A COLLABORATIVE SHIELD

Fraudsters often operate across multiple banks to hide their tracks. While individual institutions only see part of the picture, privacy laws like GDPR and CCPA often limit the direct sharing of raw customer data.[41, 42] Federated Learning (FL) offers a solution by enabling distributed training across institutions while keeping sensitive information decentralized and private.[43, 44] In a federated learning setup, multiple institutions collaboratively train a global fraud model. Each participant keeps their raw data local and only shares encrypted model parameters or "gradients" with a central server for aggregation.[41, 43]

- 1) Mitigating Overfitting: By learning from a wider distribution of fraud patterns across different banks, FL models generalize better and are more likely to catch emerging threats like synthetic ID fraud and mule accounts.[41, 45]
- 2) Reduced Data Leakage: Since raw transaction data never leaves the institutional boundary, the risk of a massive data breach is minimized, enhancing compliance with data localization and privacy laws.[43, 44]
- 3) Efficiency: Using frameworks like Flower or PySyft on platforms like Amazon SageMaker, institutions can scale their fraud detection capabilities without the need for massive data migration or centralization.[41, 43]

This "shared shield" approach allows for the identification of multi-bank scam networks that would be invisible to any single player in the financial ecosystem.[42]

VI. THE ADVERSARIAL BATTLEFIELD: ATTACKS AND DEFENCES

As AI models become the primary line of defence, they also become prime targets for "Adversarial Machine Learning" (AML) attacks. Fraudsters use malicious techniques to manipulate ML models into making incorrect predictions or bypassing security protocols.[46, 47, 48]

A. Attack Vectors in Fraud Detection

Adversarial attacks exploit the unique ways AI models process information, often targeting their underlying mathematical vulnerabilities.[46, 47]

- 1) Evasion Attacks: These are "production-time" attacks where input data is subtly tweaked to fool a trained model. For example, a fraudster might split a large transaction into smaller increments or slightly alter their digital fingerprint to bypass an existing classifier.[33, 46, 48]
- 2) Poisoning Attacks: During the training phase, an adversary injects mislabelled or malicious data into the training set. This "corrupts" the model's learning process, potentially creating a "backdoor" that allows fraudulent activity to pass undetected in the future.[46, 47, 48]
- 3) Model Extraction Attacks: By repeatedly querying a bank's fraud detection API, an attacker can build a replica of the model's decision boundaries. This "stolen" model can then be used to test and refine evasion tactics offline.[46, 47, 48]
- 4) Inference-Related Attacks: These exploit the model's output to reconstruct sensitive training data (model inversion) or determine if a specific individual's data was used in the training set (membership inference).[46, 47]

B. Advanced Defensive Maneuvers

In response to these threats, financial institutions adopt proactive and layered defences. The NIST AI Risk Management Framework (RMF) advises organizations to conduct systematic "adversarial testing" to reveal vulnerabilities before deployment. [48]

- 1) Adversarial Training: This involves augmenting the training dataset with adversarial examples, teaching the model to recognize and resist deceptive inputs.[46, 49]
- 2) Defensive Distillation: A secondary "student" model is trained to mimic the "softened" output probabilities of a "teacher" model. This creates a smoother decision boundary that is more resistant to gradient-based attacks.[48, 49]
- 3) Input Preprocessing and Feature Squeezing: By reducing the precision of input data or filtering out high-frequency noise, institutions can neutralize the subtle perturbations used in evasion attacks.[49, 50]
- 4) Anomaly Monitoring and AI Security Posture Management (AI-SPM): Continuous real-time monitoring identifies when a model's behaviour deviates from expected performance, signalling an on-going attack and allowing for human intervention. [48]

VII. GLOBAL REGULATORY FRAMEWORKS: THE EU AI ACT AND BEYOND

The regulatory landscape for AI in finance is increasingly complex, with the EU AI Act serving as a global benchmark for risk-based regulation.[51, 52, 53]. The EU AI Act classifies AI systems based on their potential to cause harm. While most high-stakes financial applications are heavily regulated, fraud detection has received a specific "carve-out" due to its importance for financial integrity. [51, 52]

TABLE IV
RISK CLASSIFICATION OF AI SYSTEMS UNDER THE EU AI ACT AND FINANCIAL APPLICATIONS

Risk Tier (EU AI Act)	Description	Financial Application
Unacceptable Risk	Prohibited systems	Social scoring; manipulative AI distorting behaviour [51]
High-Risk	Regulated systems requiring conformity assessments	Creditworthiness evaluation and credit scoring [51, 52, 54]
Exempt from High-Risk	Specific law enforcement and security tasks	AI systems used for the purpose of detecting financial fraud [53, 54, 55]
Limited/Transparency Risk	Systems requiring disclosure	Chatbots; emotion recognition; biometric categorization [51, 52]


VIII. FUTURE DIRECTIONS: TOWARD AUTONOMOUS INTEGRITY

Future development should focus on integrating neurosymbolic reasoning with deep learning, while agentic AI can automate repetitive compliance activities under human oversight. Federated learning and adversarial defenses can further support privacy and resilience. These directions point toward fraud-management systems that are accurate, explainable, privacy-aware, and capable of continuous adaptation. As global losses from payment fraud are expected to reach \$43.47 billion by 2028, the demand for accountable, transparent, and resilient AI systems has never been more urgent.[9, 38] The ongoing development of federated ecosystems, combined with robust adversarial defences and trauma-informed victim protocols, represents a holistic approach to safeguarding the global financial system. The ultimate goal is a state of "autonomous integrity," where financial networks can detect, report, and recover from fraudulent interventions with minimal human intervention, ensuring a secure and seamless experience for legitimate users worldwide.

REFERENCES

- [1] Year-over-Year Developments in Financial Fraud Detection via Deep Learning: A Systematic Literature Review - arXiv, <https://arxiv.org/html/2502.00201v2>
- [2] Year-over-Year Developments in Financial Fraud Detection via Deep Learning: A Systematic Literature Review - arXiv, <https://arxiv.org/html/2502.00201v1>
- [3] Comparative Analysis of Machine Learning and Deep Learning Models for Card Fraud Detection - Diva-portal.org, <https://www.diva-portal.org/smash/get/diva2:1884902/FULLTEXT01.pdf>
- [4] A Review of Artificial Intelligence for Financial Fraud Detection - MDPI, <https://www.mdpi.com/2076-3417/16/4/1931>
- [5] AI Fraud Detection in Banking 2026 Guide - Emburse, <https://www.emburse.com/resources/ai-fraud-detection-in-banking>
- [6] AI-DRIVEN ASSET TRACING IN STOLEN ASSET RECOVERY - International Anti-Corruption Academy (IACA), https://www.iaca.int/media/attachments/2026/02/02/iacare_final.pdf
- [7] Behavioral Biometrics in Fraud Detection - Innovify, <https://innovify.com/insights/behavioral-biometrics-fraud-detection/>
- [8] Scripta Technica: Journal of Engineering and Applied Technology Explainable AI (XAI) Analysis Using SHAP for Credit Card Fraud, <https://scriptaintechnica.com/scripta-technica/article/download/602/474>
- [9] Financial Fraud Detection Using Explainable AI and Stacking ... - arXiv, <https://arxiv.org/pdf/2505.10050>
- [10] Financial Statement Fraud Detection Through an Integrated Machine Learning and Explainable AI Framework - MDPI, <https://www.mdpi.com/1911-8074/19/1/13>
- [11] Financial Fraud Detection Using Explainable AI and Stacking Ensemble Methods - arXiv, <https://arxiv.org/abs/2505.10050>
- [12] An Introduction to Machine Learning Methods for Fraud Detection - MDPI, <https://www.mdpi.com/2076-3417/15/21/11787>
- [13] Year-over-Year Developments in Financial Fraud Detection via Deep Learning - arXiv, <https://arxiv.org/pdf/2502.00201>
- [14] (PDF) Deep Learning vs. Traditional Machine Learning Models in Detecting Claims Fraud: A Comparative Performance Analysis - ResearchGate, https://www.researchgate.net/publication/401047069_Deep_Learning_vs_Traditional_Machine_Learning_Models_in_Detecting_Claims_Fraud_A_Comparative_Performance_Analysis
- [15] Behavior-Level Models for Fintech Fraud Detection - KITRUM, <https://kitrum.com/blog/ai-fraud-detection-fintech/>
- [16] Internet fraud transaction detection based on temporal-aware heterogeneous graph oversampling and attention fusion network - PMC, <https://pmc.ncbi.nlm.nih.gov/articles/PMC12680363/>
- [17] GNN Feature Engineering for Credit Card Fraud Detection: A Comprehensive Research Framework - APIC, <https://apic.id/jurnal/index.php/jsc/article/download/546/220>
- [18] What Is Behavioral Biometrics: How Does It Work Against Fraud - Feedzai, <https://www.feedzai.com/blog/behavioral-biometrics-next-generation-fraud-prevention/>
- [19] (PDF) Financial Fraud Prevention with AI-Based Behavioral Biometrics - ResearchGate, https://www.researchgate.net/publication/387172274_Financial_Fraud_Prevention_with_AI-Based_Behavioral_Biometrics
- [20] Behavioural Biometrics and Device Intelligence in Fraud Strategies - Experian, <https://www.experian.co.uk/blogs/latest-thinking/guide/behavioural-biometrics-device-intelligence/>
- [21] AI agents for fraud detection and investigation: how they work and ..., <https://www.unit21.ai/blog/ai-agents-for-fraud-detection-and-investigation-how-they-work-and-what-to-evaluate>

- [22] Can AI Write Your Suspicious Activity Reports? AI-Generated SAR Narratives, <https://www.sanctionsanner.com/blog/can-ai-write-your-suspicious-activity-reports-ai-generated-sar-narratives-1340>
- [23] How AI Agents for Financial Crime Run the Full Compliance Lifecycle - Blog | Unit21, <https://www.unit21.ai/blog/how-ai-agents-for-financial-crime-run-the-full-compliance-lifecycle>
- [24] NICE Actimize's Transaction Monitoring Software - STAR, <https://www.niceactimize.com/anti-money-laundering/suspicious-transaction-activity-reporting>
- [25] Achieve 60% Higher SAR Filing Rate With Automation + Gen AI, <https://www.automationanywhere.com/resources/use-case/financial-services/improve-aml-process-efficiency>
- [26] Understanding Suspicious Activity Reports: A Detailed Guide - Unit21, <https://www.unit21.ai/fraud-aml-dictionary/suspicious-activity-report>
- [27] AI-Powered SAR Narratives & Automated Filing Software | Hawk, <https://hawk.ai/platform/sar-filing>
- [28] 12 Key Questions to Guide Victim Notification Protocols - Sexual Assault Kit Initiative, <https://www.sakitta.org/resources/docs/10391R3SAKI12KeyQsVctmNtfctnPrctl.pdf>
- [29] Security & Fraud Protection | Northwest Bank, <https://www.northwest.bank/security-center/>
- [30] Asset Recovery Guidance and Best Practices - FATF, <http://www.fatf-gafi.org/content/dam/fatf-gafi/reports/Asset-Recovery-Guidance-Best-Practices.pdf.coredownload.pdf>
- [31] Practices.pdf.coredownload.pdf
- [32] AI and OSINT for Financial Investigations and Asset Recovery - unodc, <https://www.unodc.org/roca/en/NEWS/2026/ai-and-osint-for-financial-investigations-and-asset-recovery.html>
- [33] Section 10.1 Suspicious Activity and Criminal Violations - FDIC, <https://www.fdic.gov/resources/supervision-and-examinations/examination-policies-manual/section10-1.pdf>
- [34] (PDF) Adversarial Attacks and Robustness of AI in Financial Fraud Detection - ResearchGate, https://www.researchgate.net/publication/395890720_Adversarial_Attacks_and_Robustness_of_AI_in_Financial_Fraud_Detection
- [35] Victim Notification Protocols for Untested Sexual Assault Kits: Survivors' and Advocates' Perspectives on Law Enforcement, <https://justiceresearch.dsplacedirect.org/server/api/core/bitstreams/8d1be63c-30df-40c4-a77f-bee13e73833f/content>
- [36] Credit Freezes and Fraud Alerts - FTC Consumer Advice, <https://consumer.ftc.gov/articles/credit-freezes-and-fraud-alerts>
- [37] Explainable AI in Finance: Meeting Stakeholder Needs | RPC, <https://rpc.cfainstitute.org/research/reports/2025/explainable-ai-in-finance>
- [38] Explainable AI for Regulatory Compliance in Financial and Healthcare Sectors: A comprehensive review - International Journal of Advances in Engineering and Management (IJAEM), https://ijaem.net/issue_dcp/Explainable%20AI%20for%20Regulatory%20Compliance%20in%20Financial%20and%20Healthcare%20Sectors%20%20A%20comprehensive%20review.pdf
- [39] Explainable AI as A Strategic Asset: Enhancing Accountability in Financial Disputes for Regulatory and Market Gains - Allied Business Academies, <https://www.abacademies.org/articles/explainable-ai-as-a-strategic-asset-enhancing-accountability-in-financial-disputes-for-regulatory-and-market-gains-17607.html>
- [40] Explainable AI (XAI) Using SHAP and LIME for Financial Fraud Detection and Credit Scoring, https://www.researchgate.net/publication/391998247_Explainable_AI_XAI_Using_SHAP_and_LIME_for_Financial_Fraud_Detection_and_Credit_Scoring
- [41] Credit Card Fraud Detection using Explainable AI Methods - Journal of Information Systems Engineering and Management, <https://jisem-journal.com/index.php/journal/article/download/3917/1724/6410>
- [42] Fraud detection empowered by federated learning with the Flower framework on Amazon SageMaker AI | Artificial Intelligence - AWS, <https://aws.amazon.com/blogs/machine-learning/fraud-detection-empowered-by-federated-learning-with-the-flower-framework-on-amazon-sagemaker-ai/>
- [43] AI-driven opportunities in financial crime prevention - Capco, <https://www.capco.com/intelligence/capco-intelligence/ai-driven-opportunities-in-financial-crime-prevention>
- [44] Federated Learning: 7 Use Cases & Examples - AIMultiple, <https://aimultiple.com/federated-learning>
- [45] The Role of Federated Learning in Improving Financial Security: A Survey - arXiv, <https://arxiv.org/html/2510.14991v1>
- [46] Federated Learning Applications: 7 Real-World Use Cases - Duality Tech, <https://dualitytech.com/blog/federated-learning-applications/>
- [47] What Are Adversarial AI Attacks on Machine Learning? - Palo Alto Networks, <https://www.paloaltonetworks.com/cyberpedia/what-are-adversarial-attacks-on-ai-machine-learning>
- [48] Adversarial AI: Understanding and Mitigating the Threat - Sysdig, <https://www.sysdig.com/learn-cloud-native/adversarial-ai-understanding-and-mitigating-the-threat>
- [49] Combating the Threat of Adversarial Machine Learning to AI-Driven Cybersecurity - ISACA, <https://www.isaca.org/resources/news-and-trends/industry-news/2025/combating-the-threat-of-adversarial-machine-learning-to-ai-driven-cybersecurity>
- [50] Adversarial Attacks and Defenses in Machine Learning - ResearchGate, https://www.researchgate.net/publication/389515260_Adversarial_Attacks_and_Defenses_in_Machine_Learning
- [51] Adversarial Attacks & Defenses in Machine Learning | by Prajun Trital | Medium, https://medium.com/@prajun_t/adversarial-attacks-defenses-in-machine-learning-088ad1993c01
- [52] High-level summary of the AI Act | EU Artificial Intelligence Act, <https://artificialintelligenceact.eu/high-level-summary/>
- [53] EU AI Act Risk Categories Explained: Why the Four-Tier Pyramid Is Wrong - Modulos AI, <https://www.modulos.ai/blog/eu-ai-act-risk-categories-explained-why-the-four-tier-pyramid-is-wrong/>
- [54] When is Artificial Intelligence High Risk? - Leaseurope, <https://www.leaseurope.org/when-artificial-intelligence-high-risk>
- [55] Annex III: High-Risk AI Systems Referred to in Article 6(2) | EU Artificial Intelligence Act, <https://artificialintelligenceact.eu/annex/3/>
- [56] EC | EU AI Act & AML/Fraud Detection: Regulatory Clarification | FinancialCrime.lu, <https://financialcrime.lu/2025/05/26/EC-AI-Act-Exemption-Fraud-AML/>



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)