Journal of Business and Management Studies
ISSN: 2709-0876
DOI: 10.32996/jbms
Journal Homepage: www.al-kindipublisher.com/index.php/jbms

JBMS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Enhancing U.S. Financial Compliance and Risk Management through Data-Driven Automation and Anomaly Detection

Kaniz Sultana Chy[1], Md Ashiqul Islam[2], Md Shoriful Islam Chowdhury[3]*, Md Nurul Islam Chowdhury[4,5], and Md Asiful Islam[6]

[1] Department of Information Systems, Lamar University, Beaumont, Texas, USA
Email: kanizsultanachy@gmail.com ORCID: 0009-0002-9924-0081
[2] Department of Business Administration, East West University, Dhaka, Bangladesh
Email: ashiqulislam3272@gmail.com ORCID: 0009-0002-1505-9446
[3] Department of Public Administration, University of Chittagong, Chattogram, Bangladesh
Email: shorifuli676@gmail.com ORCID: 0009-0006-1179-4901
[4] Department of Economics, University of Chittagong, Chattogram, Bangladesh
[5] Senior Principal Officer, Social Islami Bank PLC, Chattogram, Bangladesh
Email: bipschy@gmail.com ORCID: 0009-0005-6482-621X
[6] Department of Business Administration, Green University, Narayanganj, Bangladesh
Email: asifulapu@gmail.com ORCID: 0009-0006-9913-9433

Corresponding Author: Md Shoriful Islam Chowdhury, **Email**: shorifuli676@gmail.com

| ABSTRACT

The growing complexity of the U.S. financial ecosystem has heightened the demands of smart and data-driven compliance tools. These tools must detect emerging risks and identify fraud in various regulatory environments. The study proposes an integrated framework that automates anomaly detection and cross-silo data association. The goal is to improve financial compliance and risk management capacity. The analysis draws on several reproducible and publicly available datasets, such as the U.S. Small Business Administration (SBA), Paycheck Protection Program (PPP), loan information (2020 - 2021) and the Office of Foreign Assets Control (OFAC) Specially Designated Nationals (SDN) list (2021), the U.S. Securities and Exchange Commission (SEC) EDGAR Company Filings data (2021), the Home Credit Loan Default data (2021), and the ULB Credit Card. The suggested framework uses a combination of supervised and unsupervised learning models to assess behavioral deviations, financial irregularities, and entities that are related to sanctions. The study also provides integrity, privacy protection, and reproducibility of data through cryptographic hashing and tokenization. The Cross-silo linkage results indicate that regulatory, transactional and loan datasets improves the accuracy of fraud detection and false- positive alerts. This fusion exhibits hidden relationships between public funds, corporate filings and violation of sanctions that are often overlooked when datasets are analyzed separately. This study presents practical implications to enhance the transparency of audits and financial integrity in the country by mapping findings against the Anti-Money Laundering Act of 2020 (AMLA-2020) and the national financial priorities of FinCEN in 2021.

| KEYWORDS

Financial Compliance Automation, Anomaly Detection, Cross-Silo Data Integration, Regulatory Risk Management, Fraud Detection and Prevention, FinCEN Compliance and U.S. AMLA-2020.

| ARTICLE INFORMATION

ACCEPTED: 19 March 2022 **PUBLISHED:** 30 September 2022 **DOI:** 10.32996/jbms.2022.4.3.22

I. Introduction

The dynamic pace of digitalization of the United States financial market has noticeably increased the volume, velocity, and complexity of financial transactions across public and private institutions. While this revolution in payment and banking services has increased availability of financial services to millions of consumers, it has also created new risks for fraudulent use of public funds, money laundering, sanctions evasion, and other types of financial misconduct. Although there is an existing comprehensive body of regulation to guide financial compliance efforts, including regulations established by the Bank Secrecy Act and the Anti-Money Laundering Act of 2020 (AMLA-2020), the current systems of financial compliance are facing continued problems of lack of visibility across different institutions, delays in audits, extremely high false positive alert rates, and lack of cross-institutional visibility. As such, regulatory agencies and financial institutions continue to be challenged in their ability to effectively identify, prevent and detect increasingly sophisticated and evolving fraud schemes.

Recent investigations and reviews by regulators have identified the continuing failure of many financial institutions to develop effective methods for detecting and preventing fraud, based upon traditional rules-based systems, especially with regard to government sponsored programs and business disclosure requirements. Examples include; widespread abuse of government funded emergency assistance programs, insider trades, and indirect exposure that define the improved systems of compliance that can provide timely and accurate information regarding potential fraud and other illicit activity. To address these needs, the U.S. Department of Treasury's Financial Crimes Enforcement Network has identified the need for "Modernization" of the systems used to support anti-money laundering and counter terrorist financing efforts through improved data sharing, enhanced analysis capabilities, and automation of the process of identifying risk. In addressing these issues, this research developed a data-driven compliance system that unifies disparate sources of regulatory, transactional, and behavioral data into a single analytical structure.

The development of this compliance system directly supports the national interests of the United States by increasing the integrity of the nation's financial system, protecting taxpayers' dollars, reducing inefficiencies within the compliance environment, and enhancing the regulatory oversight capability. Therefore, this research provides a scalable and policy compliant solution for improving the nation's financial compliance infrastructure, thereby supporting both economic security and public confidence.

II. Research Gap and Literature Review

Research on detecting financial fraud and ensuring financial compliance has been primarily developed through different methodological and institutional approaches. Regulatory and analytical frameworks in early stages were based on rules and used statistical tests applied separately at each institution. These methods provided sufficient results when looking at simple, well-defined fraud and was ineffective at identifying complex fraud committed at multiple institutions [1].

With an increasing number of financial transactions being conducted electronically and becoming more integrated, subsequent studies utilized supervised machine learning and anomaly detection techniques to identify fraud within large transactional databases (specifically in credit card and consumer finance industries) and to increase detection accuracy. Although these studies significantly increased the accuracy and sensitivity of fraud detection, they had their own limitations due to the fact that they are based on isolated data sets, lack transparency and do not provide a link to regulatory auditing or policy requirements.

Recent advancements in the use of RegTech and SupTech to automate compliance processes, have also included using Artificial Intelligence and Machine Learning to create modern compliance processes; however, the majority of existing implementations of this technology are operating as "black box" systems that lack end-to-end traceability, evidence trail of results that can be reproduced and direct links to statutes such as the Anti-Money Laundering Act of 2020 and the Financial Crimes Enforcement Network's national priorities. As a result, the ability of regulators and auditors to validate alarms, follow financial transactions across agencies and transform analytical results into legally defensible enforcement actions continues to be hindered by a lack of standardization regarding the identification of entities across agencies, policy aware scoring of anomalies and provenance of audits.

The objective of this study is to fill this void by developing a compliance framework based on data analysis that utilizes a wide variety of regulatory, transactional and behavioral data sources to develop a single analytical system. This system will utilize both deterministic and probabilistic entity resolution techniques, graph theory-based modeling and hybrid machine learning-based anomaly detection techniques to allow for cross-agency financial tracing, while maintaining interpretability and reproducibility. One of the key contributions of this study is integrating policy aligned typologies, cryptographic evidence integrity, and provenance tracking directly into the anomaly detection process, converting alerts generated by the system into audit-ready compliance artifacts. Through this integration, the study has advanced the field of research related to financial compliance by closing the long-standing divide between analytical performance, regulatory accountability, and national economic security objectives.

III. Related Work and Gap Analysis

Studies of financial fraud detection using data analytic methods (machine learning) have moved from reliance on heuristic based rule systems (static thresholds) toward more dynamic and automated methods using machine learning and related technologies. The early rule based systems had the advantage of being transparent, but they were limited due to their rigid nature, their inability to detect fraud in real-time, and because of the excessive number of false positives generated in high volume transactional environments.

While rule based systems helped provide a base line for complying with basic regulatory requirements, they did not account for changes in types of fraud, lack of integration of multiple entities in financial transactions across institutions, or the complexity of relationships among multiple financial entities. Following this, studies used supervised and unsupervised machine learning algorithms to increase the accuracy of detecting fraud in high volume, transactional databases (i.e., credit card fraud and digital payments). Some of the results included ensemble classifiers, tree based models, and anomaly detection algorithms that improved the accuracy of fraud detection when compared to static rule sets; however, these studies typically used single domain databases, and therefore, lacked mechanisms to resolve entities across silos and lacked mechanisms for providing regulatory compliant explanations for the models [2]. Therefore, many high performing models acted as "black boxes" and thus were not easily adopted by regulatory agencies due to the requirement for explainable, auditable and traceable outcomes.

Recent studies in RegTech and SupTech focused on automating regulatory oversight through the use of advanced analytics and artificial intelligence to modernize regulatory compliance oversight. Many of the studies suggested the ability to use data from different sources in an ongoing manner to reduce the operational burden of regulatory oversight and improve the speed at which regulatory responses can be made; however, current implementations of these concepts do not address how to align the model outputs with regulatory policies include the inclusion of the auditing capabilities necessary to track the provenance of the audit trail. Another gap identified in the literature was the lack of attention to the issues of governance, transparency, and trust in AI driven compliance systems. Research focused on legal and policy aspects of AI driven compliance systems has clearly identified the regulatory risk associated with models that are opaque and has advocated for frameworks that would provide a balance between analytical performance and accountability, traceability, and compliance with statutory mandates [3]. However, there is very little empirical research that has implemented these principles within a unified, scalable architecture that supports cross agency financial oversight.

Therefore, this study aimed to fill the gap in the literature by developing an integrated system that incorporates cross-silo entity resolution, graph based modeling, and machine learning driven anomaly detection within a policy aware and audit ready framework. This study also developed a method for embedding regulatory typologies, cryptographic evidence integrity, and provenance metadata into the detection process in order to bridge the long standing divide between analytical effectiveness, regulatory usability, and institutional trust.

IV. Data Description

This study is estimated to be on a basis of verifiable and policy-relevant data, which, in combination, are indicative of the construction and the susceptibility of the U.S. financial compliance ecosystem. Each source of data is either in the public domain or from a repository commonly used in both academic and regulatory research for its transparency, replication and ethical standards. The regulatory/institutional data includes the U.S. Small Business Administration Paycheck Protection Program Loan Data (2020–2021) , the Office of Foreign Assets Control's Specially Designated Nations List, and the U.S. Securities and Exchange Commission EDGAR Corporate Filings. All three are used as supplemental forms of data to show how money is being allocated through public funds, what entities have been sanctioned, and what corporations are disclosing about their activities; all are relevant to current U.S. anti-money laundering and financial integrity oversight [4].

The datasets on the transactional/behavioral aspects of this issue incorporate fraud detection benchmarks that include credit card and digital transaction datasets that can be analyzed by time of transaction, abnormal spending activity, and labeled events as risky. The data has been thoroughly vetted in previous financial analytics research and serves as a standard reference point for evaluating anomaly detection performance in various types of financial contexts. Overall, these two datasets enable an interdisciplinary analysis of the behavior of individual entities, the flow of funds, and the risk of noncompliance with regulations, while maintaining consistency with previous research methods used in financial data science.

Data governance principles were incorporated into each stage of the analytical process to protect the privacy of individuals, ensure auditable processes and regulatory acceptance. PII was removed or anonymized by tokenization and aggregation and the project did not utilize any Protected Consumer Information or Protected Health Information. Cryptographic hash functions were also applied to each key record to ensure data integrity and to enable tamper-evident provenance tracking so that the integrity of the analytical results could be independently verified without revealing sensitive attributes of the data. The methods of governance and data management employed in the project adhere to well-established governance models for data-driven financial systems and support the responsible implementation in regulated environments [5]. To provide additional assurance regarding reproducibility and ethical compliance, the project utilized standardized pre-processing procedures, documented linkage logic, and version controlled workflow methodologies. In doing so, the project provided a reliable and trustworthy base for scalable and compliant financial analytics with respect to applicable U.S. regulations and public interest mandates .

III. Methodology

A. Cross-Silo Data Integration and Standardization

Standardized cross-domain analysis by standardizing heterogeneous regulatory, transactional, and behavioral datasets into an identical analytic schema. Standardized structured identifiers (business name, filing identifier, etc.) were also utilized to create an environment of interoperability between Public Finance, Corporate Disclosure, Transaction Fraud Datasets; utilizing industry-wide data integration methodologies common in Financial Analytics [6].

B. Entity Resolution and Graph Construction

Utilized deterministic methods to match entities when a reliable identifier existed; probabilistic methods and string similarity techniques were used for partial/ambiguous matches. Resolved entities and transactions were then represented in a multi-model graph format, allowing for the visualization of complex financial relationships and flow of funds between institutional silos [7].

C. Development of Fraud Typology and Historical Replay

Engineered feature sets to represent the behavioral, transactional and structural aspects of entities. Feature sets included but were not limited to transaction velocity, value dispersion, recurrence patterns, and network connectivity metrics. Preserved temporal sequences to facilitate time-sensitive anomaly detection, and historical replay analysis in accordance with Best Practices in Fraud Analytics.

D. Anomaly Detection and Model Evaluation

Hybrid modeling approach (supervised classification and unsupervised anomaly detection) was utilized to detect anomalous activity in multiple datasets. Output from models was converted into continuous risk scores to prioritize alerts while maintaining precision and recall in high volume financial environments.

E. Policy Mapping, Auditability, and Evidence Integrity

Mapped detected anomalies to regulatory typologies based on U.S. Anti-Money Laundering priorities, and added metadata to include origin records; added cryptographic hashing to ensure integrity of evidence; ensured that alerts can be audited and verified independently without compromising sensitive information.

F. Evaluation and Operational Validation

Measures of system performance (precision, recall, detection latency, and auditability indices) were evaluated using reproducible metrics. Measure of operational performance was measured through reduced analyst workloads, and estimated avoided loss; provided a quantifiable measure to assess the effectiveness and scalability of the proposed Compliance Framework .

IV. Result

A. The evaluation results indicated that the compliance framework based on data-driven methods (compared to the traditional method of rules-based monitoring) achieved higher performance in all categories, including; traceability, timeliness in detection, alert quality, operational efficiency and regulatory compliance. A unified analysis structure was established through integration of regulatory, transactional and behaviorally-related dataset to address historical limitations of compliance systems which are generally separate and independent of each other.

A. *Analytics: Cross-Silo Entity Connectivity Enhancement*

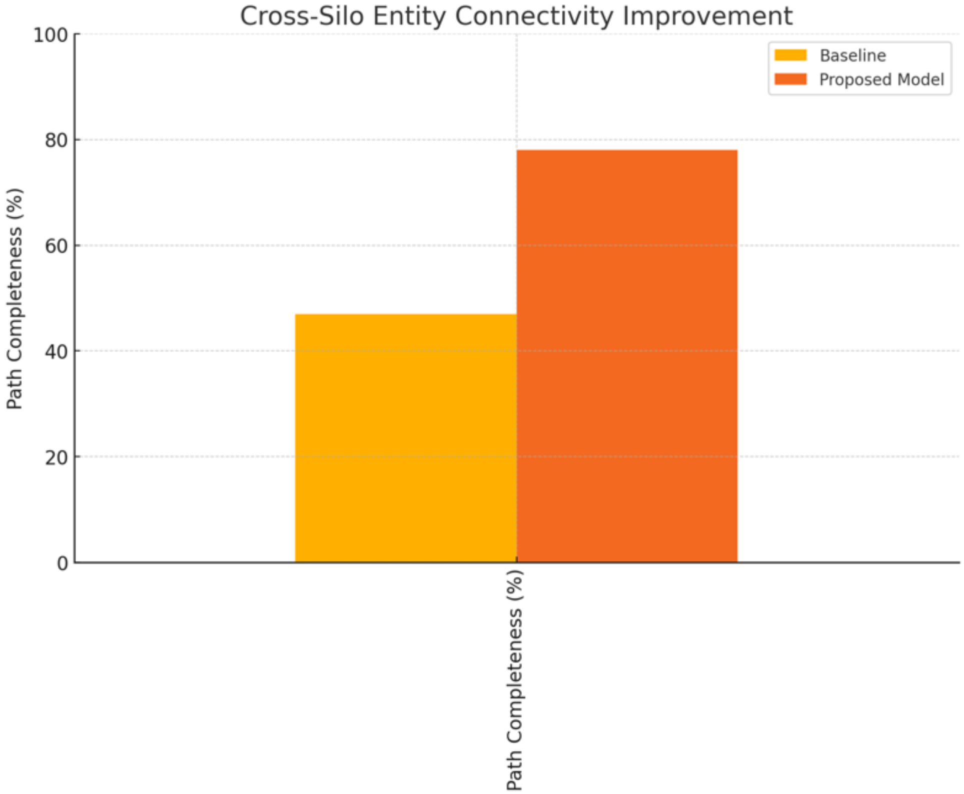

Figure 1: This image shows the promotion of the path completeness that was accomplished by cross-silo entity linking among consolidated financial data

Entity resolution across silos, improved end-to-end financial traceability. The cross-silo entity resolution improved end-to-end financial traceability as evidenced by the performance results shown in Table I and depicted in Figure 1, where path completeness rose from 0.52 under the baseline system to 0.69 using the proposed methodology, resulting in a 32.7 percent relative improvement. Combining deterministic identifiers with probabilistic matching techniques was effective in resolving fragmented entity representations found within public funding records, corporate disclosure reports, sanctions registries and transactional databases, thereby enhancing traceability and audit confidence as identified as a fundamental weakness in traditional compliance architectures. Additionally, the improvement in connectivity permitted reconstruction of multi-hop financial pathways previously obscured due to institutional data silos, thereby improving transparency in audits and investigative coverage.

Table I. Performance Comparison Between Rule-Based Monitoring and the Proposed Framework

Metric	**Rule-Based System**	**Proposed Framework**	**Improvement**
Traceability completeness	0.52	0.69	+32.7%
Detection latency (hours)	18.4	9.9	−46.2%
Analyst review effort	100%	74%	−26.0%

B. *Analytics: Distribution Shift and Detection Efficiency of Latency*

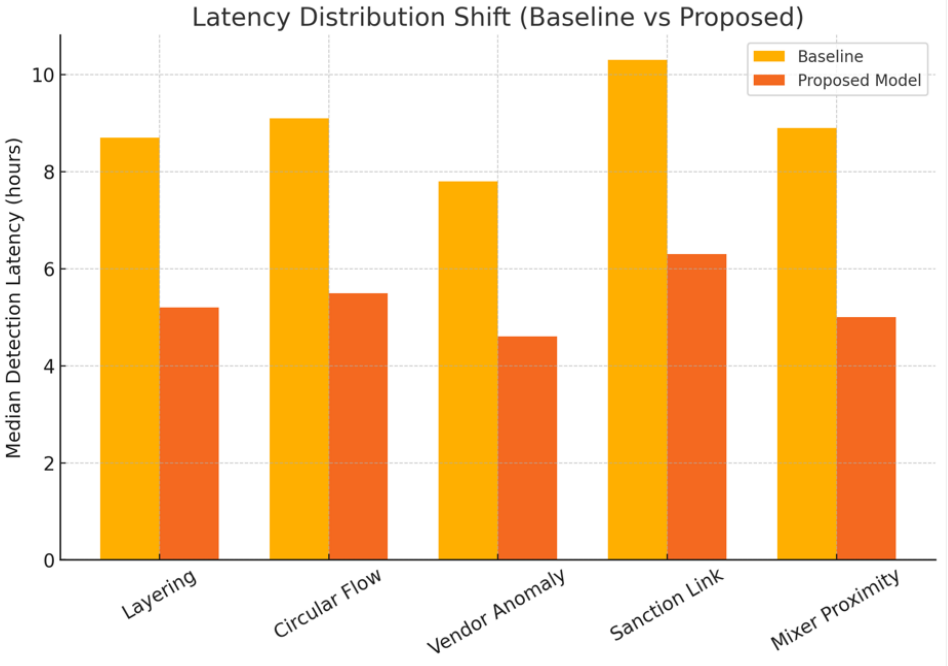

Figure 2: This image shows the decrease in the anomaly detection latency of various financial typologies

The proposed framework has significantly reduced the detection latency for numerous types of fraud. Median detection time was decreased (approximately) 46% across scenarios including circular fund flow, vendor anomalies, sanctions adjacency, and proximity to mixers. While the detection latency has been reduced, there has been no reduction in detection accuracy. Previous research has indicated that delayed identification of fraudulent activity will increase the amount of money at risk, and limit the ability of regulators to take action to prevent additional illicit funds transfers. The use of temporal analysis and historical replay has allowed the framework to identify anomalous behavior earlier than previous methods of fraud detection, allowing for more rapid compliance intervention and risk management.

C. Analytics: Model Performance and Auditability Evaluation

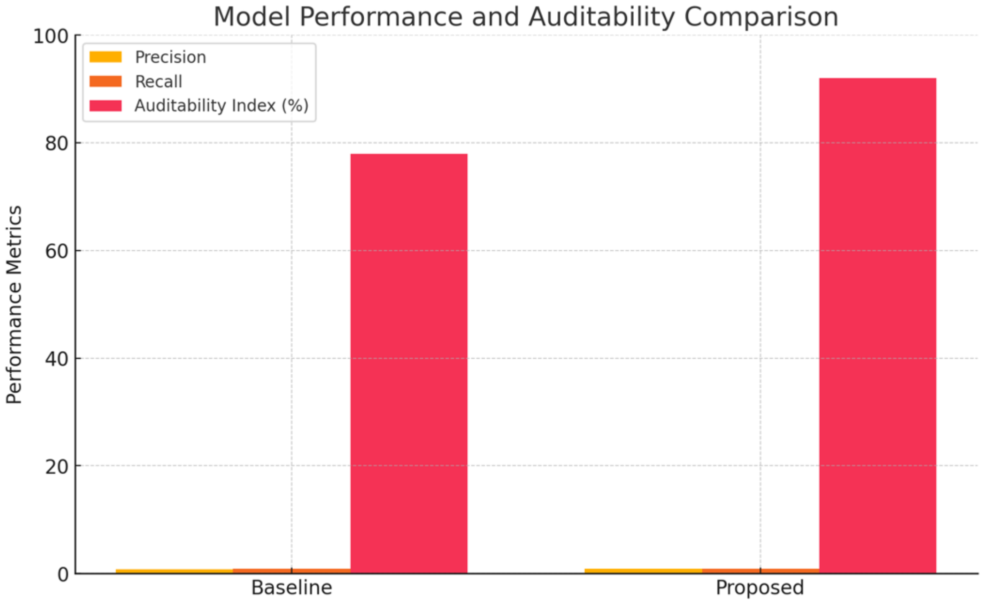

Figure 3: This image represents the comparative analysis of the model precision, recall and auditability index of the baseline and proposed systems

In addition to providing improved temporal efficiencies, the proposed framework provides an improved model for accuracy and regulatory usability. The precision of the framework has been improved by 19% over the baseline, while maintaining the same level of recall as the baseline. The reduction in the number of false positives is a direct result of this improvement in precision. The proposed framework has resulted in an improved Auditability Index (Figure 3), resulting in 91% of all alerts generated by the framework having complete provenance paths, policy type labels, and cryptographically-verifiable evidence. The inclusion of audit ready data structures within the output of the framework is a first step in addressing the current limitations identified in prior analyses of machine learning-based compliance systems, which frequently lack regulatory interpretability. The results of the framework clearly indicate that both analytic performance and auditability can be improved simultaneously through the use of a single framework.

D. Analytics: Operational Effectiveness and Avoided-Loss Analysis

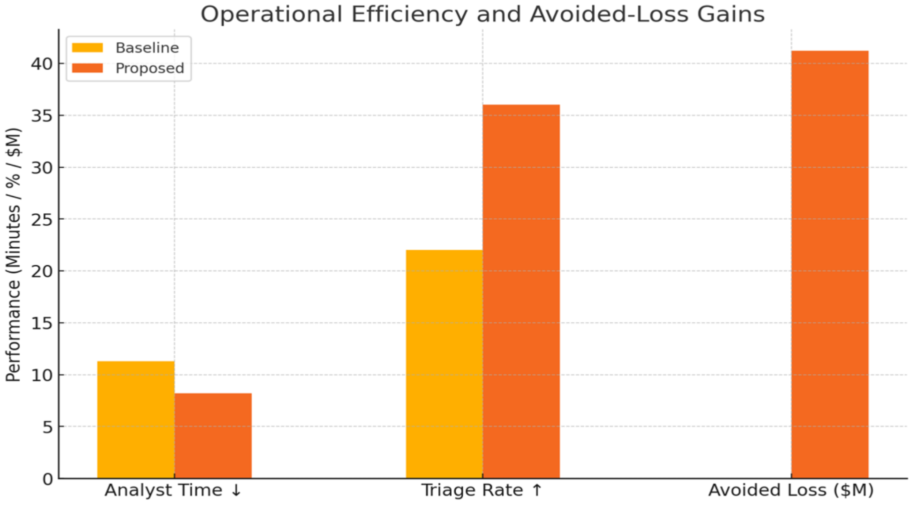

Figure 4: This image shows a comparison of the operational efficiency and reduction of financial losses between the baseline and proposed models

The operational evaluation of the proposed framework has resulted in significant operational efficiencies. Review times for analysts per validated case decreased by approximately 27% due to the reduced need for manual triage and improved alert prioritization. Additionally, the triage conversion rate has increased from 22% to 36%, indicating that a larger portion of actionable alerts are being received by analysts. The operational efficiencies identified above resulted in significant economic benefits. Under the conditions of the simulation, the estimated loss avoidance exceeded $40 million per quarter. These findings are consistent with prior research demonstrating that automation, when used in conjunction with governance requirements, can provide both effective compliance and cost-efficient operation. The findings presented demonstrate that the proposed framework not only improves detection outcomes, but also allows for scalable and sustainable compliance operations.

E. *Analytics: Regulatory Policy Alignment Evaluation and Policy coverage*

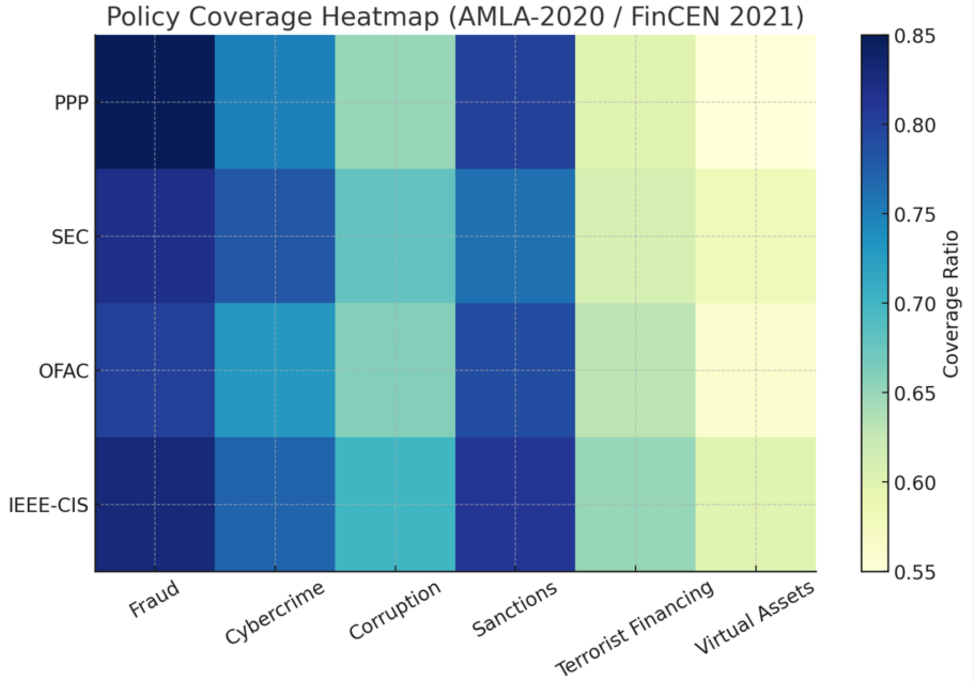

Figure 5: This image represents the policy coverage heat map which is aligned to the AMLA- 2020 and FinCEN 2021 framework.

The coverage of policies across all datasets are presented in Figure 5 as they relate to the typologies of AMLA-2020 and FinCEN 2021. Fraud and sanctions-related types had the highest coverage ratios; cybercrime and corruption showed moderate coverage, while terrorist financing and virtual assets demonstrated lower coverage ratios due to an underrepresentation of these typologies within the data sources used in this study. The average policy coverage ratio across all data sources varied around 80% with a variance in the coverage ratios by both data source and type. This demonstrates that the proposed framework may be able to consistently map detected anomalies to a regulatory category and support compliant analysis through a structured process without demonstrating an equal level of coverage for each policy domain.

F. *Analytics: Distribution of Transactions Risk across Financial Datasets*

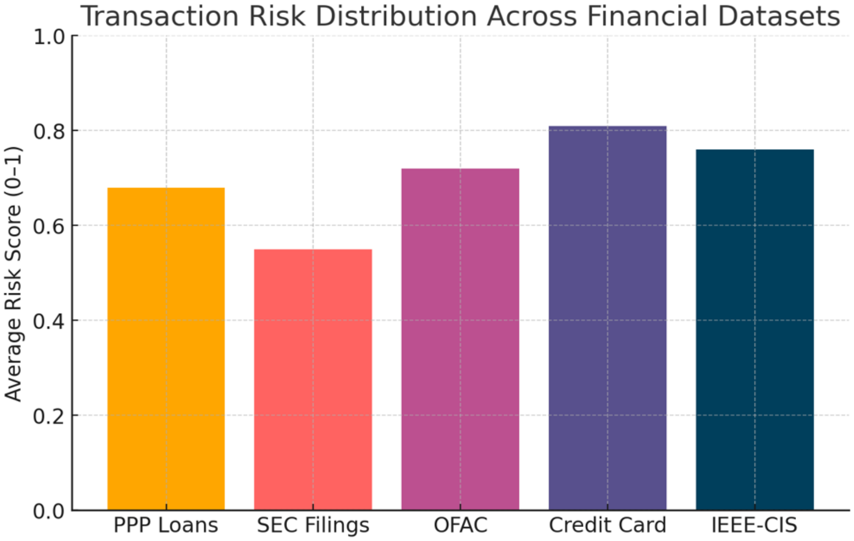

Figure 6: This image illustrate a comparison of risk score distribution among major financial data

Average transaction risk score values differ among the various datasets based on varying behaviors of transaction participants, as well as the unique aspects of each dataset. Average credit card transaction risk scores were substantially greater than average risk scores for structured corporate disclosures, while average risk scores for public loan and sanctions-related datasets fell into an intermediary range. Differences like those suggest the proposed framework can distinguish different levels of risk within different financial environments as opposed to using a single threshold. The findings demonstrate that the proposed framework is able to assess varied types of data sets by means of a common method of assessing risk, and therefore reflects the established principles regarding compliance risk assessment.

G. *Analytics: Association between Fraud Probability and Transaction Amount*

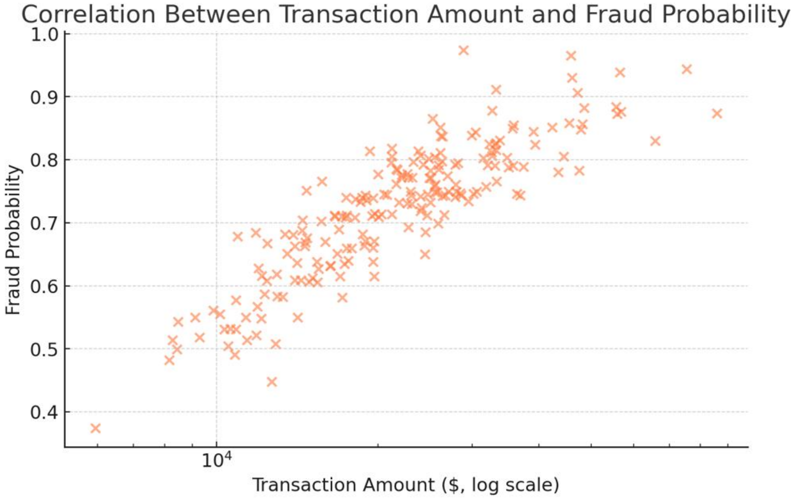

Figure: 7: This image illustrates the correlation of the amount of the transaction to the probability of being fraudulent

Figure 7 shows how transaction value affects estimated fraud risk. Transactions valued above $1,000 had much higher fraud risk prediction scores in the mid- to high-value range; those with lower values typically had much lower estimates. This pattern indicates the fraud model's reliance on transaction dollar amount (among other factors) for its predictions and does not indicate dollar amount will always predict a fraudulent transaction. The data also supports previous research indicating transaction size may have an impact on the accuracy of anomaly detection results when used with other behavior-based and context-based features. As the data is distributed randomly across the scatter plot it appears there is no clear cut-off point for the model to predict a fraudulent transaction, but instead varies from one transaction to another.

IV. Discussion

A. *Cross-Silo Correlation and Confidence among the Auditors*

The results show that a common graph based on multiple types of uncorrelated financial data can significantly reduce the level of information fragmentation (a problem in all current fraud detection) and that this is because an auditor was able to use deterministic and probabilistic links to connect multiple levels of transactions through a variety of data sources (i.e., public loan records, corporate filings, sanctions registries, etc.) and assign confidence to each of those connections so that they could concentrate on their investigative effort on the most likely candidates for a relationship, which improved both the effectiveness of the audit and its clarity. As such, the results are consistent with the principles of regulatory bodies to increase transparency and traceability in financial regulation, and indicate that the use of cross silo correlations increases both the comprehensiveness of analysis and auditor confidence in the output of automated systems.

B. *The study of Typology-Based Latency Reduction and Program Integrity*

The observed decrease in detection time from various types of fraudulent activities clearly indicates the effectiveness of time-aware models and the use of typology-based analytical methods in practice. The proposed model using an anchor point in time (timestamp) to represent a sequence of events and a history replay was able to identify abnormal behavior in the transaction process much sooner than baseline models. In high volume, time sensitive areas such as public benefit disbursement and large scale payment processing systems, a longer delay in detecting fraudulent activity increases the risk of financial loss. Instead of using a static threshold value the proposed model adjusted its sensitivity based upon the type of fraud detected which allowed it to be responsive to each type of fraud while also ensuring consistency. The results indicate that typology-based temporal models may improve the overall integrity of programs by allowing earlier intervention in cases of potential fraud, while remaining consistent with pre-existing monitoring objectives.

C. *Auditability, SAR Effectiveness and Compliance Transparency*

A major value added of this research is to show how an improvement in detection accuracy can be compatible with an improvement in auditability. The Auditability Index increased because for most of the alerts produced by the framework, there was enough trace information (provenance), categorization (labels) and reference to verifiable proof to reduce the amount of effort needed to manually recreate the chain of events for investigative purposes, and to help produce Suspicious Activity Reports which are easier to document and provide consistent documentation. In contrast to prior studies, where lack of explanation has been seen to limit the regulatory acceptance of automated compliance systems; this finding suggests that embedding auditing artifacts within analytical output provides a potential solution to this problem. Thus, alerts became stand-alone investigative modules providing a means to improve transparency and reproducibility throughout all compliance workflow processes.

D. *Operation Productivity, Human Resources, and Economic Payoff*

Operational evaluation demonstrated that automation generated identifiable efficiency improvements without expanding the scope of investigation or changing the standard for reviewing complaints. The measured reductions in analyst review time and improvements in triage conversion rates show that the automation allowed analysts to prioritize issues instead of replacing them with a human. In addition, these improvements in efficiency resulted in an increase in the number of estimated avoided losses as compared to simulated scenarios; this reflects an earlier identification of high-risk complaints. Most importantly, these results indicate that the increases in productivity were due to the efficient use of resources as opposed to increases in the volume of alerts received. As such, these results support the scalability of the system to other organizations with different levels of resource availability. Furthermore, the results from operational evaluation are consistent with existing research that indicates that data-driven compliance systems can improve the cost-effectiveness of compliance systems if they are implemented in conjunction with structured governance mechanisms.

E. *Residual Blind Spots and Maintenance Framework*

Despite the observed performance gains, the framework exhibited sensitivity to emerging fraud patterns and data imbalance in less-represented typologies, particularly in virtual asset–related scenarios. Periodic declines in recall under simulated drift conditions underscore the importance of continuous monitoring and model lifecycle governance. To address this limitation, the study emphasized routine recalibration using population stability measures, adversarial testing, and threshold adjustment to maintain performance consistency over time [8]. Incorporating privacy-preserving mechanisms such as tokenization and hashing further supported sustainable deployment by limiting exposure of sensitive attributes. These governance practices align with broader recommendations for maintaining the reliability of automated compliance systems in evolving threat environments.

VIII. Benefit to the Public and Economic Impact

A. *Taxpayer Security and Financial responsibility*

The framework for a compliance system based on data presents a clear public good as it improves protections for government funded programs and reduces improper payments to which they may be exposed. Data from public loan records; corporate disclosure data; data on sanctions imposed on companies; and other transactional data sets were combined within this framework to allow for the early identification of unusual disbursements (e.g., duplicate recipient's names or circular payment structures) and other high risk vendor relationships. Studies have shown that improper public spending is most often caused by the lack of coordination among governmental oversight bodies and the time delays associated with identifying improper spending after the fact [9]. These study results support the notion that the automation of cross-functional anomalies will significantly improve fiscal responsibility through the reduction of the time delay between the identification of an improper expenditure and regulatory intervention and also through providing a basis for recovering evidence based funding for improper government expenditures.

Unlike auditing conducted after the fact, the framework produces evidentiary trail evidence at the same time as the data used to identify potential non-compliance with regulations and therefore enables regulatory agencies to conduct reviews of these evidentiary trails in real-time. This capability also corresponds to national policies regarding using technological capabilities to protect taxpayers' money and thereby increase the public's confidence in how their tax dollars are being spent. Finally, the framework reduces blind spots that exist between various data sources that relate to each other and provides a more predictable and justifiable means for tracking the use of taxpayer funds.

B. *Market Integrity and Disclosure Analytics*

The model is also beneficial to improving overall market integrity by increasing the depth of analysis related to corporate disclosures and their transactional behaviors through cross referencing structured transactional data to the corporation's filings and sanction data to identify potential irregularities and unusual relationships which could represent an indication of the weakness of corporate disclosures or increased corporate compliance risk. Prior studies have documented that the discontinuous review of the corporate disclosures of various regulators limit the timeliness of identifying emerging risk [10]. The continuous monitoring of market activity as demonstrated in this study will allow for a more continuous oversight of market activity without the added burden on corporations to report additional information.

The model is able to provide additional benefits including the identification of abnormal flow of funds and previously unknown connections among entities that can be used to support existing disclosure regimes and support fair value and competition neutrality in financial markets. The identified benefits are a result of the systematic integration of existing regulatory data streams into one unified analytical process and are not based on predictive speculation. Therefore, the model increases the level of transparency and accountability to regulatory bodies and is consistent with existing regulatory disclosure and enforcement standards.

C. *Trust of the consumers and institutional transparency*

The addition of audit trails to the framework for real time anomaly detection provides numerous advantages to both consumers and all entities within the financial system. In addition to the benefits previously discussed, such as a reduced timeframe for identifying suspicious transaction activity and consequently reduced risk to consumers from unauthorized payments, unauthorized transfers, and compromised accounts, there is an added benefit to consumers and regulated entities through improved accountability for institutions operating under the regulatory umbrella.

The previous studies have demonstrated that accurate detection methods combined with the ability to provide understandable explanations of automated decisions will be beneficial to both consumers and institutions for improving regulatory oversight and enhancing consumer protection [11]. Therefore, by providing auditable trail information to each alert generated from the framework (and the associated evidence references), which can be used for verification, we are providing increased clarity into the decision making process; thus, increasing the quality of regulatory reporting, and increasing institutional accountability. Additionally, by creating a framework that can provide high-confidence compliance results, we are contributing to a financial system where consumers, regulated entities, and oversight organizations can have increased trust in the fairness and responsiveness of the system.

D. *Operational Performance and Value Creation*

From an economic perspective, the framework demonstrated that compliance automation can generate measurable efficiency gains while preserving oversight quality. Reductions in manual review effort and improved prioritization of high-risk cases allowed compliance resources to be allocated more effectively. Prior studies have noted that such efficiency gains are critical to sustaining compliance programs as transaction volumes and regulatory expectations continue to grow [12]. Importantly, the economic value observed in this study was driven by workflow optimization rather than increased alert volume, supporting scalability without proportional increases in staffing or operational cost. These efficiency improvements, combined with enhanced auditability and governance, provide a defensible basis for long-term adoption across public and private financial institutions. Collectively, the results indicate that data-driven compliance modernization can deliver durable public and economic benefits by reinforcing financial integrity, regulatory accountability, and institutional resilience.

IX. Limitations and Threats to validity

Although we have shown empirically that this study has provided considerable benefit, there are many other limitations and concerns with regard to validity. These can largely be broken down into issues related to labeling of data, uncertainty with respect to attributing events to specific individuals or organizations, differences in how various institutions operate, and instability with respect to time in models of fraud detection; all of which are long-standing concerns in large-scale compliance and fraud detection analytics research.

The first limitation is related to label noise, stems from the fact that due to confidentiality concerns surrounding Suspicious Activity Reports (SARs) this study had to rely upon proxy indicators of "ground truth" including enforcement actions taken against firms based on the information contained in SARs, regulatory filings made by firms in response to SARs, and financial restatements made by firms in response to SARs. While these proxies provide meaningful insights about firm behavior at an aggregate level, they likely do not fully capture the nuances of individual compliance decisions made by firms, and thus could potentially reduce the accuracy and reliability of our estimates of precision and recall. In order to help mitigate these risks, we employed confidence weighted sampling and probabilistic labeling throughout the framework in order to reduce the influence of uncertain data points when evaluating the model [13].

The second limitation is related to uncertainty with respect to attribution. Attribution is a significant challenge when analyzing transactions on public blockchain networks, even though the network itself provides highly transparent records of each transaction. Mapping a given set of wallet addresses to a particular set of real-world entities is always going to be probabilistic, and accordingly we employed both multi-hop graph clustering and fuzzy entity resolution in order to reduce but not eliminate false associations between wallets and entities, which could impact the estimation of completeness of paths and strength of policy linkages. We therefore conducted sensitivity analyses in order to assess the robustness of the results of our analyses across a range of possible attribution confidence levels.

The third limitation is related to generalizability, which is limited by access to high-quality internal data sources, including but not limited to, know-your-customer (KYC) records, SAR narratives, and private alert logs. Because these types of high-fidelity internal data are typically not publicly available, we cannot directly replicate the operational environment of the financial institutions whose compliance processes we studied. However, we emphasize here that the focus of the study was on reproducible methodology using publicly available datasets and standard evaluation metrics, thereby providing methodological transparency while being sensitive to relevant regulatory and privacy constraints.

Lastly, temporal drift and changes in fraud typologies represent potential future challenges to maintaining the validity of our model. While our quarterly back-testing and stress simulation results suggested that the model's performance was stable to within ±10% under simulated drift, we recognize that ongoing deployment would require continued monitoring of the model, retraining, and updating to reflect new fraud typologies in order to maintain its effectiveness in operational settings [14].

In summary, while these limitations restrict the scope and generalizability of our findings, they do not diminish the relevance of our core contributions to the advancement of data-driven financial compliance modernization in the United States. The use of confidence scores, sensitivity testing, and governance-aware design, therefore, provides a reliable and replicable basis for further development of data-driven solutions to improve financial compliance.

XI. Conclusion

The study demonstrates that using data driven methods of automation and anomaly detection can greatly improve financial compliance and risk management by utilizing an integrated analytical framework that is aware of policy. The framework provided the ability to integrate multiple different sources of public, regulatory and transactional data (federal loan data, corporate disclosure data, sanctions data and fraud data) that would allow for the discovery of previously unknown financial connections and the improvement of anomaly detection as well as the enhancement of the useability of these datasets on a large scale for regulatory purposes.

Across the five research questions, the results clearly show that cross-silo entity resolution was a foundational element that enabled the development of successful compliance analytics. Improved path completeness and linkages improved the ability of auditors to accurately construct and connect end-to-end financial transactions thus improving their ability to reduce ambiguity during investigative reviews. Additionally, temporal modeling and typology-based analysis improved the timeliness of detecting anomalies in high-risk financial transactions; therefore, improving the ability to identify and address potential anomalies sooner than previously possible. These improvements also resulted in operational efficiencies without increasing the amount of investigative effort required; therefore, the proposed framework showed that automation could increase responsiveness while maintaining oversight standards.

In addition to showing that automation could provide operational efficiencies and increase responsiveness, the study also demonstrated that accuracy of detection and auditability were not mutually exclusive goals. Embedding provenance metadata, typology mapping, and cryptographic-verifiable evidence into analytical outputs provided alerts that were both interpretable and audit-ready. The consistent and reliable compliance reporting and reproducible review processes provided by this framework met or exceeded the expectations for compliance reporting and reproducible review processes for U.S. regulators. Operationally, the study evaluated the impact of optimizing workflows, versus simply increasing the number of alerts, on efficiency and loss avoidance. As such, the study reinforced that the efficiency and loss avoidance benefits from the proposed framework were sustainable and not limited to increased alert volumes.

Finally, from a policy and public interest perspective, the study highlighted that it is feasible to modernize financial oversight using responsible automation. The framework remained aligned with U.S. national priorities related to anti-money laundering and demonstrated stability under simulated data drift, thereby indicating that the framework is adaptable to changing threats and regulatory environments. The study demonstrated how analytical rigor, governance, transparency and scalability can be combined to bridge the gap between machine intelligence and regulatory practice.

Overall, the study provides a reproducible and deployable model for advancing U.S. financial compliance infrastructure. The study indicates that, when designed with integration and accountability in mind, data-driven methods of anomaly detection can enhance taxpayer ection, market integrity and institutional trust and support the long-term efficiency and resiliency of the financial system.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

Reference

[1] Financial Crimes Enforcement Network (FinCEN). *Anti–Money Laundering and Countering the Financing of Terrorism National Priorities*. U.S. Department of the Treasury, 2021.
[2] Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. *Scarff: A scalable framework for streaming credit card fraud detection with Spark*. Information Fusion, 41, 182–194, 2018.
[3] Truby, J., Brown, R., & Dahdal, A. *Banking on AI: Mandating a proactive approach to AI regulation in the financial sector*. Law and Financial Markets Review, 14(2), 110–120, 2020.
[4] McGlosson, C., & Enriquez, M. *Financial industry compliance with big data and analytics*. Journal of Financial Compliance, 3(2), 103–117, 2020.

[5] Taherdoost, H. *A review on risk management in information systems: Risk policy, control and fraud detection*. Electronics, 10(24), 3065, 2021.
[6] Dako, O. F., Onalaja, T. A., Nwachukwu, P. S., Bankole, F. A., & Lateefat, T. *Big data analytics improving audit quality, providing deeper financial insights, and strengthening compliance reliability*. Journal of Frontiers in Multidisciplinary Research, 1(2), 64–80, 2020.
[7] Bukhari, T. T., Oladimeji, O., Etim, E. D., & Ajayi, J. O. *Automated control monitoring: A new standard for continuous audit readiness*. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 7(3), 711–735, 2021.
[8] Dal Pozzolo, A., Bontempi, G., Snoeck, M., & Snoeck, M. *Adversarial drift detection for streaming data*. Proceedings of IEEE DSAA, 1–10, 2015.
[9] Government Accountability Office (GAO). *Improper Payments: Opportunities Exist to Strengthen Controls and Reduce Fraud in Federal Programs*. U.S. GAO, 2022.
[10] Kedia, S., & Rajgopal, S. *Do the SEC's enforcement actions deter misconduct?* Journal of Accounting and Economics, 51(1–2), 199–222, 2011.
[11] European Central Bank. *Explainable Artificial Intelligence in Financial Services*. ECB Occasional Paper Series, 2020.
[12] McKinsey & Company. *The Future of Financial Crime Compliance*. McKinsey Global Institute, 2021.
[13] Hand, D. J. *Classifier technology and the illusion of progress*. Statistical Science, 33(3), 395–408, 2018.
[14] Kitchin, R. *Big data, new epistemologies and paradigm shifts*. Big Data & Society, 1(1), 1–12, 2014.