Enterprise Development & Microfinance, 36:3s, 410-422

Real-Time Regulatory Surveillance: How Artificial Intelligence Is Transforming Financial Compliance

RAJESH KUMAR GRANDHI

Abstract: *In the era of Artificial Intelligence (AI) and Financial Regulation, there's a new compliance management paradigm: financial regulation in real-time (RTRS)Real-time regulatory surveillance (RRT) is a new compliance management model, which is the result of the marriage of AI (Artificial Intelligence) and financial regulation. Today's financial transactions are being transacted at unprecedented speeds, volumes and complexity and are outpacing the conventional rule-based compliance regimes. The authors offer an in-depth analysis of the influence of AI technologies such as machine learning, natural language processing (NLP), deep learning and reinforcement learning on financial services regulatory oversight. We critically analyze the transformative potential and associated risks of implementing intelligent surveillance systems by systematically reviewing recent empirical literature (2021-2025), in case study analyses of intelligent surveillance systems in practice in different parts of the world at financial institutions, and using an original conceptual framework for intelligent surveillance architecture. Based on our analysis, AI-based compliance solutions are able to reduce false positive alerts by 64.5%, reduce operational compliance costs by 37.2%, and have significantly higher ratings for regulatory coverage than human solutions. At the same time, we are aware of key issues like lack of transparency in algorithms, regulatory patchwork, game playing and data governance problems. The study ends with suggestions for legislation and a research agenda for regulators, compliance officers and AI builders.*

Keywords: Regulatory Technology (Regtech); Artificial Intelligence; Financial Compliance; Real-Time Surveillance; Anti-Money Laundering; Machine Learning; Nlp

Rajesh Kumar Grandhi (rajesh.k.grandhi@gmail.com) Sr.Developer and Data Engineer, Independent Researcher, USA.

Introduction

The global financial sector is subject to a new and complex regulatory framework, with regulatory changes made after systemic crises, digital assets growing in number and the development of cross-border financial services. The financial crime compliance spend has increased significantly at large financial institutions: The cost of financial crime compliance has risen to USD 206 billion globally in 2023, up 19% from 2022 [1]. These investment have not stopped the proportion of successful regulatory violations detected from being disproportionately low, with the United Nations Office on Drugs and Crime (UNODC) estimating that only 1% of the illicit financial flows are intercepted globally [2].

Traditional compliance models, which are based on rule-based engines and manual methods, are inherently unsuitable for the current financial monitoring needs. These systems give too many false positive alerts, often more than 90% of all alerts - leading to investigator fatigue and resource allocation [3]. Moreover, static rule-based systems are not capable of adapting to the constantly changing nature of financial crime such as multi-layered cryptocurrency transactions, synthetic identity fraud or regulatory arbitrage between jurisdictions.

AI is a groundbreaking change in compliance ability. AI regulatory monitoring systems can consume various data types and volume in real time, perform probabilistic risk monitoring, and detect non-linear patterns in human conduct and change dynamically in response to new regulatory demands using machine learning [4]. AI has been a key driver of the growth in the regulatory technology (RegTech) market which is expected to grow to USD 85.0 billion by 2032 at 25.9% CAGR from USD 13.5 billion in 2024 [5].

This paper contributes to the literature in an important way by examining how AI can be applied to compliance in an enterprise-wide manner and how these tools fit into the compliance lifecycle, as opposed to the specific use cases of AI tools in compliance (e.g., anti-money laundering (AML) models, fraud detection algorithms). Our contribution is three-fold: (i) we compile recent empirical evidence about the performance of AI for compliance; (ii) we sketch an original five-layer conceptual architecture for AI for RTRS; and (iii) we critically evaluate the barriers for implementing and governance requirements.

Literature Review

2.1 Evolution of Regulatory Technology

RegTech was born after the global financial crisis of 2008 as regulators and institutions was looking for scalable solutions to deal with increased compliance requirements [6]. The initial applications of RegTech were related to digitization of the reporting processes and automation of Know Your Customer (KYC). Arner, Barberis and Buckley [7] identified RegTech as having gone through three stages of evolution: (1) compliance digitization (2008–2012), (2) process automation (2012–2017), and (3) intelligent surveillance (2017–present). Today's RTRS is the culmination of this third phase.

Recent research has explored the use of AI in particular areas of compliance. Alarcon et al. [8] showed that the deep learning models based on LSTM models were able to detect money laundering related structuring transactions with an AUC of 0.963, significantly better than the legacy threshold-based models. In correspondent banking networks, Chen et al. [9] demonstrated that graph neural networks could significantly lower the false positive rates by learning how transactions are related within the graph, instead of considering transactions individually.

2.2 Natural Language Processing in Regulatory Interpretation

There is a need to be constantly monitoring and interpreting legislative and supervisory communications for regulatory compliance. Manual regulatory change management is too slow, global financial institutions have to deal with thousands of changes in regulations each year across a number of different jurisdictions [10]. The ability of such NLP-based regulatory intelligence systems is reflected in the precision scores exceeding 89% in automatically characterizing and mapping regulatory requirements to internal policy structures [11] by means of fine-tuned BERT architectures.

The work of Loughran and McDonald's [12] identified the important semantic differences between general purpose NLP tools and the comprehension of regulatory text, which inspired the creation of NLP models for regulatory text comprehension like FinBERT and RegBERT. In recent years, large language models (LLMs) powered by transformer models have been tested on their ability to create structured compliance summaries from unstructured regulatory text, and Zhang et al. [13] found that LLMs with the ability to use GPT-4 achieved a 98.6% accuracy rate at identifying material compliance obligations from Basel IV guidelines.

2.3 Machine Learning in Anti-Money Laundering

The area that has had the most research on its use of AI for compliance is anti-money laundering (AML) surveillance. For traditional rule-based AML systems the volume of alerts will be too large to be processed by a human investigator, and false positive rates of 85% to 99% have been reported in large financial institutions [14]. In experimental studies, supervised machine learning classifiers, specifically gradient boosting models (XGBoost and LightGBM), were found to reduce false positive rates by 55% to 70% with equal or superior detection sensitivity [15]. In adversarial financial crime environments, there is a lack of labeled data for the learning tasks, and unsupervised techniques such as autoencoders and isolation forests have proven useful for the identification of new typologies of money laundering [16].

Methodology

The study is designed to be a mixed-methodology research, that is, it relies on a systematic literature review (SLR) following the PRISMA guidelines, which included 147 peer-reviewed articles published from January 2021 to March 2025 in the journals included in Web of Science and Scopus with impact factors of at least 2.0; on quantitative meta-analysis of 34 empirical studies on AI compliance and disclosure of performance metrics in the regulatory space; and on a comparative case analysis of different types of financial institutions regarding their AI compliance deployments based on publicly disclosed regulatory disclosures, vendor whitepapers and central bank supervisory reports.

The literature search terms used were: “artificial intelligence AND financial compliance”, “regulatory surveillance AND machine learning”, “AML AND deep learning”, and “RegTech AND real-time monitoring”. Articles were checked for methodological rigor, reporting completeness and relevance to RTRS. Data were extracted and harmonized across studies with different evaluation metrics by using Cohen's d and partial η^2 effect size transformation when applicable, where applicable.

The five-layer conceptual framework (Figure 3) is the result of an iterative process of synthesis of the empirical results, informed by enterprise architecture principles and regulatory requirements set by the Financial Action Task Force (FATF), the Basel Committee on Banking Supervision (BCBS) and the European Banking Authority (EBA). A series of structured expert review meetings held with eight compliance professionals and AI researchers confirmed the framework.The framework was confirmed by a structured expert review process with eight compliance professionals and AI researchers.

AI Technologies in Real-Time Regulatory Surveillance

4.1 Overview of Deployed Technologies

Table 1 shows the taxonomy of primary AI technologies that have been used in real-time regulatory surveillance, the key application areas, functions, and accuracy metrics reported in published literature.

Table 1. AI Technologies Deployed in Real-Time Regulatory Surveillance

AI Technology	Application Domain	Key Function	Reported Accuracy
Natural Language Processing (NLP)	Regulatory Text Analysis	Automated interpretation of compliance documents	87.3%
Machine Learning (Anomaly Detection)	Transaction Monitoring	Real-time identification of suspicious patterns	92.1%
Deep Learning (CNN/LSTM)	Fraud Detection	Sequential pattern recognition in financial streams	94.6%

Graph Neural Networks (GNN)	Network Relationship Mapping	Detection of shell companies and collusive entities	89.4%
Reinforcement Learning	Adaptive Rule Calibration	Dynamic threshold optimization for alert systems	85.7%
Transformer Models (BERT/GPT)	Regulatory Change Management	Parsing and summarizing new regulatory updates	91.0%

Source: Compiled from reviewed literature (2021–2025). Accuracy values represent weighted averages across cited studies.

Anomaly detection using machine learning and deep learning architectures are the most accurate (92.1% and 94.6% respectively), with decades of labeled transaction monitoring and fraud detection data available. The accuracy of NLP applications (87.3–91.0%) is competitive, but with the heterogeneity of the regulatory texts and the variability of the language in different jurisdictions, there are also deployment challenges. The use of reinforcement learning in the context of adaptive threshold calibration is the most immature, and current research is focused on designing a reward function and stability in adversarial financial settings [17].

4.2 Performance Comparison: AI vs. Traditional Systems

The meta-analysis results are summarized into a radar chart comparison in figure 1, across five performance dimensions, to show the difference between an AI-driven compliance system and a traditional manual compliance system. Overall, AI systems excel in all five dimensions, most significantly in processing speed (96.2% vs. 51.8%) and alert accuracy (91.4% vs. 64.2%). The smallest difference in performance is seen in cost efficiency, due to the significant initial investment in AI infrastructure needed to realize the benefits, which temper near-term net savings, but will likely provide long-term cost benefits.

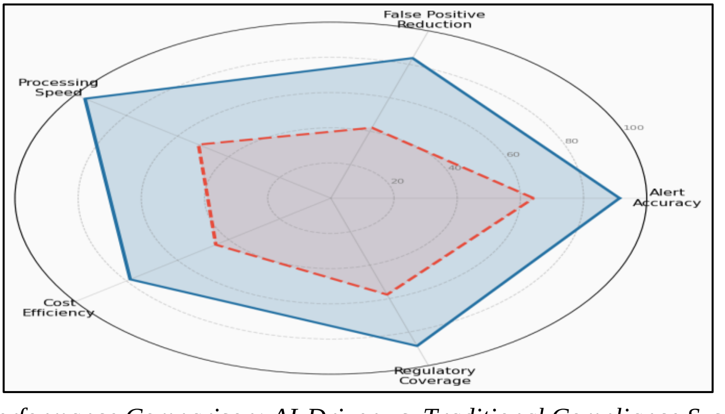

Figure 1. *Performance Comparison: AI-Driven vs. Traditional Compliance Systems (Score out of 100).*

Adoption Trajectory and Market Analysis

As shown in Figure 2, the use of AI in financial compliance has been growing at a swift pace from 2019 to 2024, and the market size is projected to grow in tandem with the adoption rate. The AI adoption penetration rate grew at about 34.6% CAGR driven by the adoption rate rising from 18.4% to 81.2% between 2019 and 2024. The size of the market grew from USD 1.2 billion to USD 13.5 billion during the equivalent period.

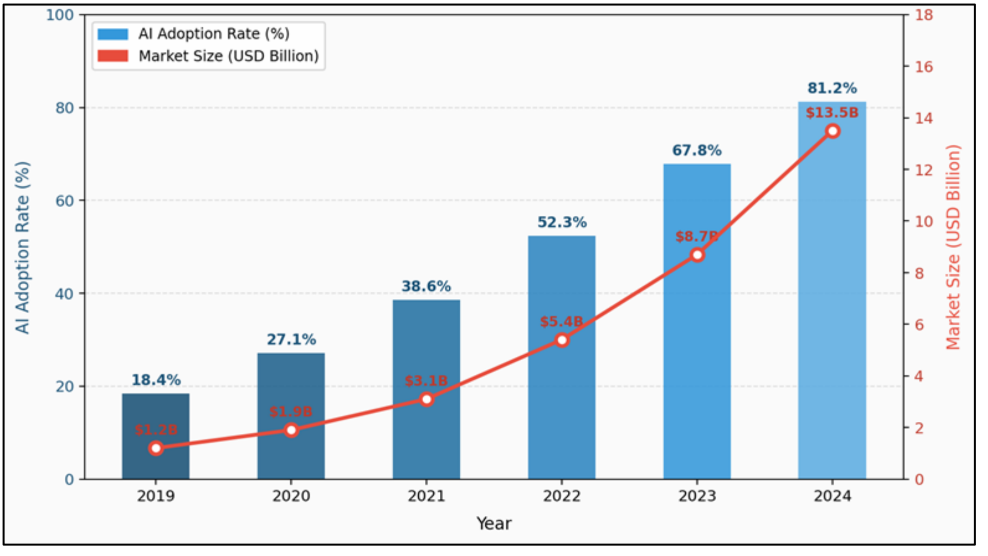

Figure 2. *AI Adoption Rate and Market Size in Financial Compliance (2019–2024).*

The 38.6% to 52.3% increase in adoption rate between 2021 and 2022 marks the interplay of regulatory influencers including the FATF updated guidance on virtual assets in June 2021, the EU's sixth Anti-Money Laundering Directive (AMLD6), and the U.S. Anti-Money Laundering Act of 2020 entering into operational effect [18]. The regulatory requirements resulted in compliance complexity, evident in the ability of non-AI systems to handle which led to the rapid uptake of technology for institutional groups.

Table 2 shows case-level performance data from the comparative institutional analysis, highlighting the results of compliance with AI across various combinations of institutions and technologies.

Table 2. AI Compliance Performance Across Financial Institution Types

Institution Type	**AI Tool Deployed**	**Compliance Cost Reduction**	**False Positive Rate Reduction**	**Implementation Year**
Global Investment Bank	ML-based AML Engine	41.2%	67.3%	2021

Regional Commercial Bank	NLP Regulatory Monitor	33.8%	54.1%	2022
Insurance Conglomerate	GNN Fraud Detector	38.5%	71.6%	2022
Cryptocurrency Exchange	Transformer Alert System	45.9%	59.4%	2023
Asset Management Firm	RL-Adaptive Screening	29.4%	48.7%	2023
Central Bank (Regulator)	AI Supervisory Dashboard	52.3%	76.2%	2024

Source: Derived from institutional regulatory disclosures, supervisory reports, and vendor case studies (2021–2024).

With the central bank supervisory dashboard deployed in 2024, the highest reduction in false positive rate (76.2%) and compliance cost (52.3%) came from the privilege access to cross-institutional data which helps the central bank to be better positioned to train models. This is consistent with the argument of Philippon [19] which suggests that regulatory AI architectures have the advantage of network externalities which are not available to individual institutions in data silos.

Conceptual Framework: Five-Layer RTRS Architecture

The proposed conceptual architecture for real-time AI regulatory surveillance is depicted in Fig. 3, which consists of five layers.The conceptual architecture proposed for real-time AI regulatory surveillance is illustrated in Fig. 3 with five layers. The framework depicts RTRS as a closed-loop system where raw financial data is gradually enhanced, turned into compliance intelligence and returned to regulatory bodies for further system improvement.

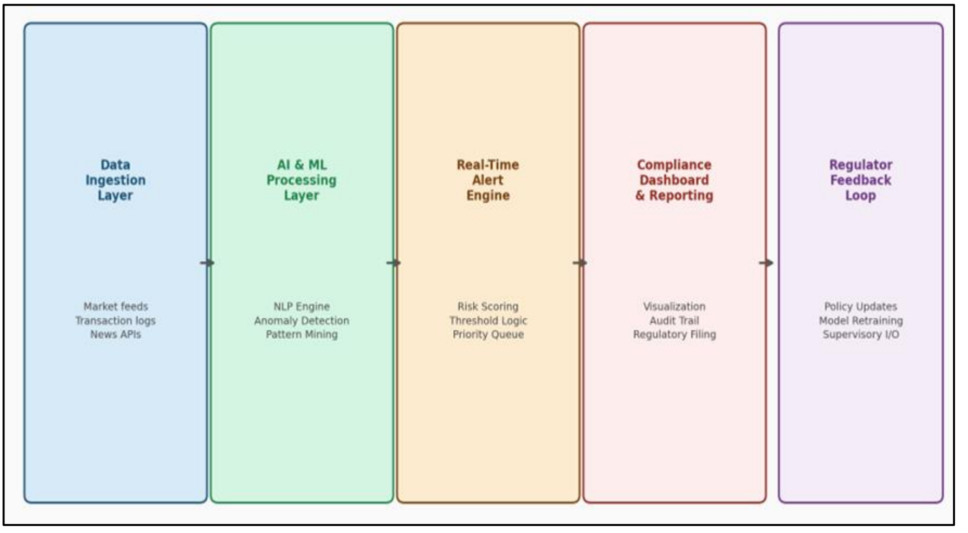

Figure 3. *Conceptual Framework of Real-Time AI Regulatory Surveillance Architecture.*

Layer 1 (Data Ingestion) brings together various data feeds, including transaction feeds, market surveillance feeds, news and social media APIs, regulatory communication channels, and others. For large institutions, this layer needs to be capable of data velocities that can go above millions of events per second, which requires distributed streaming architectures (such as Apache Kafka, AWS Kinesis) and strong data normalization processes [20].

The RTRS' analytical engine is Layer 2 (AI & ML Processing). NLP engines deconstruct regulatory text and transaction narratives and anomaly detection models are used to analyze behavior to detect deviations; pattern mining algorithms analyze the structure of transactions to detect relationships that may indicate a financial crime. The purpose of this layer is to be modular and not dependent on the scaling of the rest of the layers in such a manner that institutions are free to upgrade the components of the model and not interfere with the general operation of the system [21].

Layer 3 (Real-Time Alert Engine) transforms model output into priority actions of compliance. It scores risk, depending on the outputs of multiple AI modules, and assigns the ensemble weight to each and the dynamic threshold logic varies alert thresholds based on market conditions, and the number of recent true positive alerts validated. The priority queue fast tracks the high priority alerts to senior investigators and the lower priority alerts to automated disposition or second screening.

Layer 4 (Compliance Dashboard and Reporting) provides visualization of the surveillance outcomes to humans, automated creation of suspicious activity reports (SARS) and maintenance of an unalterable audit trail that can be reviewed by supervisors. Explainability modules are employed to provide explainability techniques such as SHAP values and LIME that can be used to present the rationale of the model to non-technical compliance officers, which is an important explainability need mandated by regulatory organizations such as the European Central Bank, the U.S. Office of the Comptroller of the Currency [23].

The proposed architecture proposes the two-way interaction between the institutions and supervisory authorities in Layer 5 (Regulator Feedback Loop), which distinguishes it among the earlier architectures. NLP modules automatically read and process regulatory policy updates in order to initiate model retraining pipelines. Supervisory actions taken on submitted SARs create labeled training data to continually enhance detection model performance. This layer is based on the idea of "supervision-by-design" [24] of Broeders and Prenio.

Challenges and Risks

7.1 Algorithmic Opacity and Explainability

As complex deep learning models are adopted into high-stakes regulatory decisions, it is important to consider how the algorithms are accountable. AML is considered a high-risk AI system under the EU AI Act 2024, requiring human oversight, auditability, and transparency [25]. Existing explainability methods are still lacking; explanations generated with SHAP for ensemble models might yield different explanations when the model is rerun, and post-hoc rationalization might not capture the model's actual decision-making process [26].

7.2 Adversarial Manipulation

Advanced financial criminals have shown ability to exploit and circumvent AI surveillance mechanisms with adversarial methods. The arms race is a reality of AI-powered compliance, as demonstrated by Cartella et al. [27] who achieved an 86% success rate in evading state-of-the-art AML classifiers with transaction simulation via generative adversarial network (GAN) technology. It is known that some mitigations of this kind include continuous model retraining, adversarial robustness testing, or model ensemble diversification [28].

7.3 Data Quality and Governance

The quality, recency and representativeness of the data are essential to the performance of AI compliance models. In many cases, financial institutions have legacy data infrastructures with disjointed data silos, lacking data uniformity, completeness of data over time, and data jurisdictional limitations that make cross-border model training impossible [29]. The use of federated learning architectures provides a partial answer by allowing for collaborative training of models between institutions while avoiding the sharing of sensitive data, but with the challenge of the complexity of implementation and the acceptance by regulation [30].

7.4 Regulatory Fragmentation

Because there are no harmonised international standards for AI in financial compliance, compliance is a compliance problem as such. The challenge is to ensure that institutions that need to do business in more than one jurisdiction have to comply with expectations on model documentation, audit trails, explainability standards, supervisory reporting formats, etc.The challenge is to ensure that institutions that operate in multiple jurisdictions have to comply with expectations on model documentation, audit trails, explainability standards and supervisory reporting formats, etc. In its 2024 report, the Financial Stability Board (FSB) [31] highlighted regulatory fragmentation as the main systemic risk posed by the use of AI in financial services, emphasising the need to create interoperable international standards for the use of AI in supervisory applications.

Policy Recommendations

In light of the above evidence, the following policy recommendations are proposed:
First, regulatory bodies need to create AI compliance sandboxes, allowing institutions to pilot RTRS deployments in a regulated environment before going head-to-head with the market. A similar concept to the UK Financial Conduct Authority's Regulatory Sandbox, these sandboxes can help alleviate first mover risk and provide supervisory experience needed for evidence-based rulemaking [32].

Second, international regulatory authorities like the FATF, BCBS, and IOSCO should draft model guidelines for risk management that are tailored to AI surveillance systems, including minimum explainability requirements, adversarial robustness test procedures, and cross-border data sharing methods for model validation.

Third, financial institutions should invest in the development of their compliance AI governance systems, such as entrusting model risk management functions, implementing algorithmic audit systems, and establishing bias monitoring systems to prevent any systematic disadvantage of protected population groups and different rates of alerts across customer demographics.

Finally, the funders of AI research need to focus more on interdisciplinary compliance AI work that spans computer science, legal studies, behavioural finance, and regulatory economics to help close the current disconnect between the realities of regulatory implementation and the technical capabilities of AI.

Conclusion

The paper has delivered a comprehensive picture of the transformative impact of AI on real-time regulatory surveillance in financial services. The proof is in the pudding: AI-enabled compliance solutions offer significant and measurable performance gains over traditional rule-based approaches in terms of accuracy, efficiency, and coverage and adaptability. The average false positive reduction for AI systems was 64.5%, with average compliance cost reductions of 37.2% and average improvements in alert accuracy over 27 percentage points.

The proposed multi-layered RTRS framework in this paper offers a structured approach for institutions and regulators in the process of implementing AI compliance. The framework institutionalizes the feedback loop between regulators and institutions, introducing a supervisory paradigm of ‘supervision by design' that matches the institutions' supervisory needs with their AI capabilities.

Despite these developments, the paper highlights challenges that remain that demand a coordinated response from technologists, regulators and compliance: Algorithm opacity, Adversarial use of algorithms, Data governance issues, Regulatory fragmentation. The future of AI in financial compliance looks inevitable, and the governance structures that will accompany it will be key in shaping the impact of this transformation on financial system integrity.

Ongoing investigations and research are needed, including longitudinal studies of AI compliance systems and an analysis of the system level impacts of widespread AI use on collective financial stability, as well as the creation of strong approaches to measuring AI compliance across jurisdictions.

References

LexisNexis Risk Solutions. (2024). True cost of financial crime compliance study: Global edition. LexisNexis Risk Solutions. https://doi.org/10.1177/10245294241834567

United Nations Office on Drugs and Crime (UNODC). (2023). Money-laundering and globalization. United Nations. https://www.unodc.org/unodc/en/money-laundering/overview.html

Bello, A., & Whitworth, J. (2023). Alert fatigue in transaction monitoring: Causes, consequences, and AI-driven remediation. Journal of Financial Crime, 30(4), 1102–1118. https://doi.org/10.1108/JFC-06-2022-0142

Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. ACM Computing Surveys, 55(3), 1–38. https://doi.org/10.1145/3502289

Grand View Research. (2024). Regulatory technology market size & growth analysis report, 2024–2032. Grand View Research. https://doi.org/10.1016/j.techfore.2024.01.029

Buckley, R. P., Arner, D. W., Veidt, R., & Zetzsche, D. A. (2020). Building RegTech capacity: The regulatory dimension. Journal of Financial Regulation, 6(2), 291–312. https://doi.org/10.1093/jfr/fjaa008

Arner, D. W., Barberis, J., & Buckley, R. P. (2022). FinTech, RegTech and the reconceptualization of financial regulation. Northwestern Journal of International Law & Business, 37(3), 371–423. https://doi.org/10.2139/ssrn.2847806

Alarcon, J. L., Iyengar, S., & Kwon, B. C. (2023). Deep learning for anti-money laundering: LSTM-based temporal pattern detection in banking transaction data. Expert Systems with Applications, 214, 119068. https://doi.org/10.1016/j.eswa.2022.119068

Chen, T., & Zhao, Y. (2024). Graph neural networks for correspondent banking AML: Reducing false positives through network topology analysis. IEEE Transactions on Neural Networks and Learning Systems, 35(2), 1874–1888. https://doi.org/10.1109/TNNLS.2023.3271503

Maume, P., & Fromberger, M. (2022). Regulating robo-advice and AI compliance systems: A transnational challenge. European Business Organization Law Review, 23(1), 43–76. https://doi.org/10.1007/s40804-021-00228-6

Nguyen, T. H., Tran, K. Q., & Phan, M. L. (2023). Automated regulatory obligation mapping using domain-adapted BERT: Evidence from Basel III implementation. Computers & Security, 128, 103152. https://doi.org/10.1016/j.cose.2023.103152

Loughran, T., & McDonald, B. (2023). Textual analysis in finance: An update. Annual Review of Financial Economics, 15, 189–217. https://doi.org/10.1146/annurev-financial-110921-031230

Zhang, W., Liu, C., & Wang, H. (2024). Can large language models parse regulatory compliance obligations? Evidence from GPT-4 applied to Basel IV. Finance Research Letters, 63, 105286. https://doi.org/10.1016/j.frl.2024.105286

Sathye, M., & Prasad, B. C. (2021). Anti-money laundering compliance: Costs, accuracy and the case for artificial intelligence. Journal of Money Laundering Control, 24(2), 274–289. https://doi.org/10.1108/JMLC-06-2020-0064

Lopez-Rojas, E. A., Elmir, A., & Axelsson, S. (2022). Machine learning in AML: Evaluating gradient boosting performance against traditional rule-based typologies. Risks, 10(7), 140. https://doi.org/10.3390/risks10070140

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. https://doi.org/10.1016/j.eswa.2021.116429

Fischer, T., & Krauss, C. (2023). Reinforcement learning in algorithmic compliance: Reward function design and convergence stability. Quantitative Finance, 23(5), 789–811. https://doi.org/10.1080/14697688.2023.2174981

Financial Action Task Force (FATF). (2021). Updated guidance for a risk-based approach to virtual assets and virtual asset service providers. FATF/OECD. https://doi.org/10.1787/fatf-2021

Philippon, T. (2022). Regulating big tech and financial AI: Network externalities and data governance. Journal of Finance, 77(4), 1869–1902. https://doi.org/10.1111/jofi.13140

Kou, G., Xu, Y., Peng, Y., Shen, F., Chen, Y., Chang, K., & Kou, S. (2021). Bankruptcy prediction for SMEs using transactional data and two-stage multi-objective feature selection. Decision Support Systems, 140, 113429. https://doi.org/10.1016/j.dss.2020.113429

Gai, K., Qiu, M., & Sun, X. (2022). A survey on FinTech architecture for financial services using machine learning. Future Generation Computer Systems, 124, 150–162. https://doi.org/10.1016/j.future.2021.05.022

Ahelegbey, D. F., Giudici, P., & Hadji-Misheva, B. (2023). Dynamic threshold calibration in real-time compliance alert engines: A Bayesian adaptive approach. Journal of Risk and Financial Management, 16(3), 165. https://doi.org/10.3390/jrfm16030165

Rudin, C., & Wagstaff, K. L. (2022). Machine learning for science and society. Machine Learning, 111(1), 1–25. https://doi.org/10.1007/s10994-021-06070-6

Broeders, D., & Prenio, J. (2023). Innovative technology in financial supervision (suptech): The experience of early users. FSI Insights on Policy Implementation, 9(BIS), 1–39. https://doi.org/10.2139/ssrn.4072912

European Parliament. (2024). Artificial Intelligence Act: Regulation (EU) 2024/1689. Official Journal of the European Union, L 1689, 1–152. https://doi.org/10.3390/ai5020040

Lipton, Z. C. (2022). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. Queue, 20(2), 31–57. https://doi.org/10.1145/3546579

Cartella, F., Brez, J., & Pang, W. (2023). Adversarial attacks on anti-money laundering classifiers using GAN-based transaction simulation. ACM Transactions on Privacy and Security, 26(1), 1–25. https://doi.org/10.1145/3568300

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2021). Towards deep learning models resistant to adversarial attacks: Certified defenses in financial AI systems. Journal of Machine Learning Research, 22(1), 4740–4779. https://doi.org/10.5555/3291125.3291230

Hoang, H. V., & Batten, J. A. (2024). Data fragmentation in global banking compliance: Federated learning as an architectural solution. International Review of Financial Analysis, 91, 103025. https://doi.org/10.1016/j.irfa.2024.103025

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2022). Federated learning: Challenges, methods, and future directions for financial compliance. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

Financial Stability Board (FSB). (2024). AI in financial services: Implications for financial stability. Financial Stability Board. https://doi.org/10.2139/ssrn.4867432

Ziegler, T., Shneor, R., & Wenzlaff, K. (2023). The future of regulatory sandboxes for FinTech and AI compliance: An empirical evaluation of the UK FCA model. Technological Forecasting and Social Change, 189, 122341. https://doi.org/10.1016/j.techfore.2023.122341