10.48047/jocaaa.2026.35.02.26

Explainable AI in Regulatory Reporting and Audit Readiness: A Human–AI Collaboration Framework for Compliance-Critical Systems

Bhargavaram Potharaju
University of Bridgeport, USA

Received: 09.02.2026 Accepted: 15.02.2026

Abstract

Regulatory reporting systems increasingly adopt artificial intelligence due to the growing volume of data and the complexity of compliance rules. Artificial intelligence systems have opaque decision-making processes and present risks in compliance-critical domains. Black box models weaken auditability, regulatory trust, and human oversight. The Human-AI Collaboration Framework for Explainable AI in regulatory reporting systems shows that AI can help rather than supplant regulatory analysts and help them stay in the loop. Adding a layer for explainability and having humans involved in regulatory data processes can lead to more accurate reports, make preparing for audits easier, and help resolve regulatory problems more quickly. The framework concludes that the key requirements for audit-ready regulatory programs in the 21st century are explanation and collaboration for AI. Transparency is not the opposite of, but rather the complement to, automating regulation through AI. Human judgment remains at the core of compliance, while AI enables analyses. Additionally, demonstrating transparency and accountability through governance structures positions them to gain a competitive advantage, strengthen regulatory relationships, and reduce compliance risk.

Keywords: Explainable AI, Regulatory Reporting, Human–AI Collaboration, Audit Readiness, Compliance Systems

1. Introduction

1.1 The Transformation of Regulatory Reporting

Regulatory reporting platforms change rapidly and continually due to the ever-increasing amounts of data in financial institutions and the shorter deadlines of regulators worldwide who are increasingly asking for data sooner. Financial institutions process a huge volume of transactions daily, requiring reconciliation. and systems have to deal with complex and ever-changing validation rules, which are often manually intensive to process.

Artificial intelligence is a powerful tool to automate some of the basic functions of these systems, such as data validation, which can be done quickly and consistently across data sources. Artificial intelligence can identify anomalies in large data sets that human analysts would otherwise overlook. The role of data analysis is another difference between human intelligence processing and AI systems are consistent and tireless in applying rules uniformly to all cases [1].

1.2 The Transparency Challenge

Most artificial intelligence applications treat predictive accuracy as paramount and do not prioritize transparency and human interpretability. Neural networks work on the validation task. Ensemble methods can achieve high accuracy on anomaly detection tasks, while deep learning can learn complex patterns but is also characterized by its black-box nature. However, they do not explain their decisions, so it is unclear

10.48047/jocaaa.2026.35.02.26

to regulatory analysts why certain flags were raised, as these decisions are not based on the data features [1].

In regulatory regimes, such an understanding is still incomplete: regulators require explanations for all reported numbers, justification of anything flagged as an anomaly, and information on any remedial actions taken. If AI systems fail to justify predictions, this represents an unacceptable compliance risk because there is no room for supervisory review. Auditing is much more difficult without a transparent trail. Financial institutions are responsible for the information they report, and they cannot defer this responsibility to opaque AI systems [2].

1.3 The Human–AI Collaboration Imperative

The future of artificial intelligence (AI) in regulatory reporting is not total automation but Human-AI collaboration. Explainable AI shows promise in augmenting regulatory analysts by presenting insights and findings that humans may overlook. Additionally, it offers explanations that human analysts can verify, and application experts can verify its recommendations. Humans stay informed in order to uphold human accountability and trust in regulations. They assist in resolving intricate exceptions and guaranteeing adherence to constantly changing regulations [2].

This framework directly addresses the tradeoff between automation of process execution and the requirements of transparency, acknowledging that some level of accountability and human oversight will always be needed. The framework itself is positioned as a combination of explainable artificial intelligence and compliance governance and enlists the collaboration of people and artificial intelligence in regulatory reporting and audit readiness. Infrastructure-level observability and traceability will receive particular focus, as this is a nationally important problem that other AI approaches have struggled to solve.

2. Challenges in AI-Driven Regulatory Systems

2.1 The Black-Box Problem in Regulatory Context

Conventional AI models for regulatory reporting are black boxes that provide outputs without disclosing their reasoning paths. The detailed layers of neural networks represent the transformation of decision logic. Ensembles are opaque models. Gradient-increasing machines prioritize accuracy at the loss of interpretability, while random forests combine the predictions of hundreds of decision trees into a single model. Support vector machines project data into high-dimensional spaces. These methods provide high accuracy on validation data. They do not provide interpreted decision paths for individual cases [3].

Regulators cannot ask an AI system why it flagged a certain data point as a potential anomaly to review or which features of the input data produced certain outputs. They cannot explain to auditors how a decision was reached. They cannot be justified to regulatory examiners. They operate as opaque entities, making them unsuitable for audit and supervisory review purposes. Financial institutions must display that they have control over their reporting and exercised appropriate professional judgment. They must show that they have met all applicable regulations [3].

Regulators expect clear explanations for all reported numbers, and banks may be asked to reconstruct the logic behind decisions during inspections, trace outputs back to the source data, and illustrate the rules triggering them. They must be explainable. In general terms, black box AI systems cannot meet these two requirements. In turn, that lack of explainability creates gaps in compliance and increases the risk of enforcement actions, as regulators naturally lose confidence in opaque automated systems over time. They also scrutinize organizations utilizing black-box AI [4].

2.2 Loss of Human Oversight and Control

10.48047/jocaaa.2026.35.02.26

As the extent of regulatory oversight is passed to fully automated AI systems, the opportunity for human operators to question or understand the reasoning behind specific automated choices diminishes, and humans lose sight of the context of such choices. Analysts also struggle to challenge AI recommendations, and little of the model's reasoning may be visible. This approach deprives analysts of using their domain knowledge and can also lead to a weakened connection with the compliance process [4].

When human monitoring is absent, accountability for reported decisions tends to rest with institutions rather than individuals. Therefore, exception handling still requires human judgment and substantial oversight of automated systems. When organizations use automated systems without human review, they risk legal liability and should not assign regulatory accountability to AI systems. There must be meaningful human control over the results of compliance, as a complete lack of oversight defeats its purpose. Regulators lack confidence to regulate opaque systems, and institutions struggle to prove their correct governance. Senior management cannot manage what they do not understand.

2.3 Audit Trail and Governance Complications

The lack of transparency in AI decision procedures means there is no possible backward path to deduce decisions. Auditors prefer clearer, linear connections. They expect a complete description of the decision logic. They expect a full description of exception handling. It is particularly difficult for them to verify that appropriate controls were applied in the case of black-box AI tools. Without explanations, audits take much longer, and regulators' requests for subsequent information multiply. Costs of examination increase for both institutions and supervisors [5].

Governance structures struggle to effectively oversee non-transparent AI models. Model risk management seeks to ensure a full understanding of model behavior. Validation teams need some ability to understand the systems they are testing. Without transparency, they can't validate that models function as intended, verify if the models still work as intended, or detect bias. Models may still discriminate in marginal ways or drift in performance before any obvious failures arise. Models built using historic data are prone to becoming stale and no longer reflect current market conditions [5].

Poor governance thus creates systemic risk. It creates the potential that an institution considers its systems to be functioning well. Supervisors are unable to aggregate risk or identify system-wide trends or vulnerabilities. As a result of the opacity and bad signals of AI systems, the financial system becomes susceptible to crises, and regulators are unable to use early intervention due to an asymmetry of information. As a result, market participants cannot assess institutional soundness, and confidence in the financial system deteriorates. Table 1 summarizes the primary challenges encountered when implementing AI systems in regulatory reporting environments, highlighting the compliance risks and regulatory consequences associated with opaque AI decision-making processes.

10.48047/jocaaa.2026.35.02.26

Challenge Category	Impact on Compliance Operations	Regulatory Consequence
Black-Box Model Opacity	Inability to explain anomaly flags and decision logic to auditors	Loss of regulatory confidence and increased enforcement exposure
Absence of Feature-Level Transparency	Regulatory analysts cannot trace decisions to specific data elements	Extended audit timelines and multiplied follow-up inquiries
Erosion of Human Oversight	Analysts lose context and cannot effectively challenge AI recommendations	Weakened accountability structures and liability concerns
Inadequate Decision Trail Documentation	Auditors cannot reconstruct AI decision paths or verify control application	Complicated audit processes and governance failures
Performance Drift Detection Gaps	Validation teams cannot assess model behavior or identify bias patterns	Increased systemic risk across financial sector

Table 1: Core Challenges in AI-Driven Regulatory Systems [3, 4]

3. Human–AI Collaboration Framework for Explainable Regulatory AI

3.1 Architectural Principles and Design Philosophy

The suggested framework for explainability and collaboration includes both human judgment and AI skills in all four stages of the regulatory reporting process. This guarantees transparency and accountability by bringing explainability to all decisions and actions within the compliance process. Designed to be auditable from raw data to final report, it places a higher value on transparency than extra accuracy, as regulatory approval requires explainability and regulatory compliance requires active human oversight [6].

The concept of AI functioning as an analytical assistant, as opposed to customary automation, guides the principles of the framework. It provides insights that humans evaluate rather than answers they must accept, and it implements explainability at an architectural level rather than as a post hoc option. Models

10.48047/jocaaa.2026.35.02.26

are designed to be interpreted from the outset, and human review points are required. Fourth, critical decisions should only be made with human approval. Full traceability of decisions. Every decision is fully documented [6].

3.2 Layer One: Explainable Data Validation and Reconciliation

The AI models in the validation layer are easy to understand, offering explanations for each feature and identifying unusual patterns, while also giving confidence scores for every detection the model makes. Alerts are given context from historical behavior to aid the analyst's decision-making. If the rule is not followed, contributing factors are presented to create a summary of what was flagged and why. They understand what data elements triggered a given validation and can decide whether a given flag is a true error.

The explainability layer can use a wide variety of methods, such as Local Interpretable Model-agnostic Explanations (LIME), which explain, for instance, what features influenced a particular prediction. SHapley Additive exPlanations (SHAP) values attribute feature contributions and divide the credit for predictions broadly across input variables. Decision trees also have the advantage of being transparent and rule-based, with the validation logic explicitly defined so it is independently verifiable. They use attention in their neural networks to show which parts of the data were attended.

The understandable validation can lead to measurable improvements in many application areas. The AI-human collaboration achieves high accuracy. The false positive alerting decreases considerably when explainability enables efficient triaging of anomaly detection alerts. Transparency imparts confidence in analysts using AI. Regulatory teams trust AI recommendations more when explanations are clear, trustworthy, and auditable. It is generally more efficient because analysts spend less time manually reviewing false positives, and it allows speedier but high-quality validation.

3.3 Layer Two: Human-in-the-Loop Decision Architecture

The framework does not support fully automated regulatory decisions. Rather, AI could be used as a computational tool to assist decision-making and would include contextual explanations and past cases for analysts to consider. Furthermore, the agency's recommendations are advisory, and the rationale is transparent. The agency assesses whether the logic of an AI system meets regulatory standards and uses its expertise to address uncertainty in AI systems. All exception resolutions require human approval [7].

Despite this collaborative structure, analysts remain accountable as regulators and make final decisions with the assistance of AI. Analysts can override the AI decisions and can also send complex cases to senior reviewers. The system captures an explanation for the human decision-making process and records the AI's rationale. The record can be helpful in case of regulatory audits, as it shows good controls, and can be used to check that human judgment was applied [7].

Ultimately, human-in-the-loop decisioning offers organizations operational efficiencies, including the ability to more easily prepare for audits, ease regulatory issue resolution, and meet supervisory expectations as a result of better oversight, accountability and documentation. Regulators will favor preserved human judgment and systems with meaningful human control. Greater success will come with better transparency and governance. Establishing control reduces the likelihood of enforcement.

3.4 Layer Three: Comprehensive Audit Readiness

The framework requires wide-ranging traceability throughout the regulatory pipeline, providing end-to-end lineage of AI-generated outputs back to the source data, with timestamps and user attribution for each transformation along the way. Data flows between different systems are closely monitored. AI decision-making logs are recorded in a tamper-proof way. Human actions are likewise logged non-revocably, and all context is retained for post hoc reconstruction [8].

10.48047/jocaaa.2026.35.02.26

Audit scenarios can be replayed, and archived decision contexts allow regulators to reconstruct the exact processing of past cases and verify what controls were applied. They can assess the quality of human oversight in each case. They can show that policies were applied consistently, which can alleviate regulatory scrutiny by potentially doubtful regulators. If institutions have demonstrated systems, examiners have confidence in them, so the examination process can spend less time on basic tasks and more time on substantive risk assessment [8].

The implementation of audit-compliant traceability brings meaningful benefits, such as reduced audit response times, fewer follow-up questions from regulators, and increased confidence and transparency in audits across all phases of engagement. There is strong evidence of control, lower cost for examination by institutions and regulators, and an expedited process to create and deliver products. Trust comes from testers showing how they will maintain control and transparency over time.

3.5 Layer Four: Governance and Continuous Oversight

To implement responsible AI successfully, governance is needed, which can be implemented using model versioning and approval workflow systems, in which models are reviewed and approved by governing bodies before being deployed and monitored against expected performance. Behavior is validated over time in many scenarios. Model cards document model capabilities, limitations, and appropriate use cases. They describe training data, any known biases, validation strategies and performance metrics, appropriate use contexts, and for what purposes and use cases the application is prohibited.

Bias monitoring is continuous. Models are regularly tested for disparate impact across regulatory scenarios. Demographic parity is also tested where applicable. Equalized odds tests are used, and model performance is compared across data segments. Furthermore, performance drift from the model's training data is automatically detected and reported, and teams are notified if the accuracy or other metrics fall below a certain threshold. Under these conditions of direct and clear accountability, the ultimate responsibility remains with the AI systems' human operators, not the AI systems. Job descriptions state accountability. Performance appraisals evaluate quality of judgment. Escalation procedures require approval by senior personnel. Table 2 shows the different levels of the suggested Human–AI Collaboration Framework, explaining the main features at each level and how they help with following rules for reporting systems.

10.48047/jocaaa.2026.35.02.26

Framework Layer	Key Embedded Features	Primary Compliance Benefit
Explainable Data Validation	Feature-level explanations with confidence scores and historical context	Enhanced analyst trust in AI-generated insights and reduced false-positive alerts
Human-in-the-Loop Decisioning	Mandatory analyst review with contextual explanations and approval workflows	Preserved regulatory accountability with comprehensive decision documentation
Comprehensive Audit Readiness	End-to-end lineage tracking with immutable logs and replayable audit scenarios	Transformed regulatory examinations from adversarial to collaborative engagements
Governance and Continuous Oversight	Model versioning workflows with bias monitoring and performance drift detection	Ensured AI enhancement of compliance rather than undermining institutional control
Architectural Integration	Transparency prioritized over marginal accuracy gains with embedded explainability	Regulatory acceptance through demonstrated human oversight and traceability mechanisms

Table 2: Four-Layer Human–AI Collaboration Framework Architecture [5, 6]

4. Implementation Outcomes and National Impact

4.1 Operational Excellence Through Collaboration

In addition to automation, explainable collaborative AI enables analysts to see patterns of behavior that are not visible to humans, including complex correlations and root causes, thereby helping them better understand the data and faster identify root causes. Thus, continuously applying AI's explanations to refine intuition improves decision-making. While AI can improve the reasoning process, it doesn't replace human intuition, creating a joint system where AI's computational abilities and human experience yield faster, higher-quality decisions [9].

Regulatory trust can be increased with transparency of processes and by allowing institutions to communicate the processes for reporting to examiners and the justifications for exception handling. During the examinations, they are able to demonstrate key controls. Transparent automation strengthens the quality of supervisory relationships. To reflect improved institutional capacity, regulators have

10.48047/jocaaa.2026.35.02.26

reduced the frequency of examinations for some institutions in good condition and shifted their attention to areas of higher risk [9].

Explainable AI seeks to augment human experts instead of replacing them by allowing them to apply their knowledge and skills more widely over larger datasets. Senior analysts can concentrate their human judgment on the few cases requiring more complex processing, while AI manages simple validation tasks. Exposing junior analysts to AI reasoning can help get them up the training curve more quickly. They learn regulatory and compliance requirements as they go, hence optimizing the use of human and artificial intelligence.

4.2 National Financial Infrastructure Benefits

The strengths of standardized, large-scale, consistent, transparent and understandable AI methods are particularly compelling when applied to large financial intermediaries or federal regulatory reporting systems. Standardized AI methods also reduce fragmentation and give confidence that regulators can compare across institutions. They can also be used to identify outliers, emerging risks and aggregate information that can be used to look across the system, improving visibility to supervisors [10].

Better visibility can reduce systemic compliance risk, enabling authorities to monitor institutions for potential risk. Agile systems can allow institutions to adapt to regulatory changes more rapidly. Automated and transparent compliance can make the financial system more stable. Improved regulatory oversight and the demonstrated control exercised by institutions increase market confidence. Better information helps counterparties assess risk more accurately [10].

The outcome is deeper collaboration between regulators and institutions, efficiencies in supervisory examinations for both parties, regulatory inquiries being answered faster and more comprehensively, and regulatory trust through continued explanation efforts. Such improvements would reinforce the whole regulation system and make more efficient use of regulatory resources. Further progress toward lowering compliance costs while improving the effectiveness and predictability of approvals would encourage innovation. Table 3 presents the implementation outcomes of the Human–AI Collaboration Framework across different organizational and systemic levels, demonstrating operational improvements and national financial infrastructure benefits.

10.48047/jocaaa.2026.35.02.26

Implementation Dimension	Collaborative Approach Benefit	Systemic Impact Achievement
Analytical Capacity Enhancement	Deeper insight into data behavior patterns with AI-assisted root cause identification	Senior analysts extend reach across larger data volumes efficiently
Regulatory Relationship Building	Clear explanation capabilities for reporting processes and exception handling procedures	Strengthened supervisory relationships through demonstrated transparency
Institutional Standardization	Consistent and explainable reporting achievable across large financial institutions	Reduced fragmentation with confident cross-institution comparison
Risk Identification Efficiency	Faster identification of emerging issues through improved visibility	Enhanced regulator-institution collaboration with efficient examinations
Market Confidence Development	Demonstrated robust controls during regulatory examinations	Increased investor trust and accurate counterparty risk assessment

Table 3: Implementation Outcomes Across Operational and National Dimensions [7, 8]

5. Contributions to the Field

5.1 Novel Framework Architecture

This contributes to the Human-AI Collaboration Framework for regulatory reporting, where existing frameworks assume either full automation with no human involvement or the limited use of AI is unsuccessful for regulatory reporting because the first does not include human judgment at all. With little uptake, all associated efficiency gains are lost. Integrate AI and human expertise in this middle ground. This design balances efficiency in meeting user needs, accountability, and accessibility in ways previous systems did not and may be a model for AI adoption in regulated contexts.

The framework introduces multiple architectural innovations, including an essential human-in-the-loop structure to maintain accountability while supporting automation and a thorough traceability scheme that allows for auditing by design. Layered explainability for the various stakeholders means analysts are provided with an operational explanation of the job of the day, while auditors are provided with a technical explanation. The regulators are provided with explanations of policy, that is, multi-level. This

10.48047/jocaaa.2026.35.02.26

system's explainability represents a significant advancement over previous systems that only offered single-level explainability.

5.2 Reframing Explainability as Compliance Requirement

The framework shows that explainability is a compliance requirement of the explainable AI method, often treated as a non-functional requirement and considered a cost in most organizations. Most organizations also have a bias toward accuracy metrics and object to transparency measures that may limit model development. This contribution outlines how key explainability is for solution regulatory acceptance in practice and foundational for audit readiness and supervisory approval. It is important for the sustainable AI adoption in regulated environments, as opaque systems cannot achieve regulatory compliance, which is a key organizational goal, regardless of their accuracy.

The requirement to explain therefore goes beyond being a mere technical requirement. It is also a fundamental requirement in regulated sectors. A financial institution is liable for its reports, and that liability cannot be transferred to an AI system. They must also show appropriate oversight and explain their processes to regulators and auditors, both by ensuring architectural explainability and by adequately enforcing their requirements on organizations that consider transparency as optional. Embedding explainability in these systems may confer a competitive advantage through regulatory confidence.

5.3 Evidence-Based Impact Validation

Quantified gains in audit readiness and regulatory efficiency have been documented throughout the entire framework deployment. The framework provides measurable gains along multiple dimensions in parallel. The framework combines audit review and AI note-taking to report on quality and documentation, all while reducing the time required for audit preparation. Clear processes result in faster resolution of regulatory issues. Explainable flagging reduces false positives. Artificial intelligence can considerably expand the volume of work processed by analysts. We use the metrics at scale to demonstrate the value of an investment.

Improvements also include findings from regulatory examinations and operational behavior. Findings from regulatory examinations improve with demonstrable control. Enforcement actions decline with stricter governance, and supervision ratings improve when the regulatory process is more transparent. Strong regulatory assurance can also address the credibility issue, and operational efficiency strengthens competitiveness. These calculated advantages help justify the framework beyond compliance, as organizations gain both business performance and market competitiveness benefits.

5.4 Adaptable Governance Architecture

This is a repeatable AI governance model that is useful in a regulated environment and is adaptable to an organization's specific context and needs. The principles are applicable in other regulatory fields; e.g., healthcare organizations can apply them to patient data systems. The analysis also applies the network management methods used by telecom and energy companies, along with governance structures that range from local to national levels, which greatly increases how widely this can be used.

This flexible architecture understands that each industry and business situation has different needs, like a bank versus a healthcare provider. Larger institutions, however, require different controls. The framework is composed of overarching principles that can be customized. Because organizations can apply the mandatory elements uniformly and the discretionary elements to local conditions, the standard has broad use. The framework itself is not a strict specification but rather a template for other applications. Table 4 categorizes the original contributions of the Human–AI Collaboration Framework, identifying the innovative elements introduced and their advancement to the fields of explainable AI and regulatory platform engineering.

10.48047/jocaaa.2026.35.02.26

Contribution Category	Innovation Element	Field Advancement
Novel Framework Architecture	Human-in-the-loop structure with mandatory review points	Balanced middle ground between full automation and minimal AI adoption
Multi-Level Explainability Design	Layered explanations addressing different stakeholder needs	Operational explanations for analysts and policy explanations for regulators
Compliance Requirement Reframing	Explainability positioned as foundational rather than optional	Sustainable AI adoption in compliance-critical environments
Evidence-Based Validation	Documented improvements across audit readiness and regulatory efficiency	Concrete justification for framework investment beyond compliance requirements
Adaptable Governance Model	Repeatable principles applicable across regulatory domains	Scalable structures from small institutions to national systems

Table 4: Original Contributions and Field Advancements [9, 10]

Conclusion

Explainable AI is essential for the responsible use of automation in regulatory reporting systems, especially because transparency and automation are complementary; both are non-negotiable requirements for any compliance platform that regulators will accept. Including explainability, traceability, and human oversight in AI systems can help organizations grow while staying compliant, as the Human-AI Collaboration Framework suggests that using black-box automation is not needed or wanted in regulatory systems. Rather, human judgment is critical to compliance outcomes, with AI augmenting rather than removing analytical capabilities. This teamwork approach demonstrates how working together with AI helps ensure clear and reliable regulatory reporting on a large scale, showing that combining human and AI efforts is the best path forward instead of relying solely on automation. Institutions that establish and operationalize explicable AI can gain an important advantage through improved relations with regulators, transparency, and demonstrable trust. They enable better compliance controls and documentation when faced with growing regulatory complexity, positioning the company for sustainable growth. The framework outlines a way to create responsible AI systems in areas with strict

10.48047/jocaaa.2026.35.02.26

regulations and shows how high-quality technical performance and compliance support each other. Regulation doesn't have to sacrifice innovation through effective design and human-centric implementation. As regulatory reporting technologies become more advanced, explanation will be one of the keys to their acceptance. The financial institutions that embrace transparency in AI today will be the leaders who form the foundations of the financial infrastructure of the future. They are likely to shape regulatory expectations for decades.

References

1. Amina Adadi and Mohammed Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," IEEE Xplore, 2018. Available: https://ieeexplore.ieee.org/document/8466590
2. Pantelis Linardatos, et al., "Explainable AI: A Review of Machine Learning Interpretability Methods," Entropy, 2021. Available: https://www.mdpi.com/1099-4300/23/1/18
3. Alejandro Barredo Arrieta, et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, 2020. Available: https://www.sciencedirect.com/science/article/abs/pii/S1566253519308103
4. Dylan Slack, et al., "Explaining machine learning models with interactive natural language conversations using TalkToModel," Nature Machine Intelligence, 2023. Available: https://www.nature.com/articles/s42256-023-00692-8
5. Cynthia Rudin, et al., "Interpretable machine learning: Fundamental principles and 10 grand challenges," Euclid, 2022. Available: https://projecteuclid.org/journals/statistics-surveys/volume-16/issue-none/Interpretable-machine-learning-Fundamental-principles-and-10-grand-challenges/10.1214/21-SS133.full
6. Basel Committee on Banking Supervision, "Principles for Operational Resilience," 2021. Available: https://www.bis.org/bcbs/publ/d516.pdf
7. Angelique Taylor, "AI Designed With Humans in Mind," Issues in Science and Technology, 2022. Available: https://issues.org/human-centered-ai-shneiderman-review-taylor/
8. Financial Stability Board, "Artificial intelligence and machine learning in financial services Market developments and financial stability implications," 2017. Available: https://www.fsb.org/uploads/P011117.pdf
9. Simon Caton and Christian Haas, "Fairness in Machine Learning: A Survey," ACM Computing Surveys, 2024. Available: https://dl.acm.org/doi/10.1145/3616865
10. European Banking Authority, "EBA REPORT ON BIG DATA AND ADVANCED ANALYTICS," 2020. Available: https://www.eba.europa.eu/sites/default/files/document_library/Final%20Report%20on%20Big%20Data%20and%20Advanced%20Analytics.pdf