Pages:1-15

Multidisciplinary Studies and Innovative Research

Volume-V, Issue-II (2024)

https://allscoperesearch.com/

Agentic Compliance-by-Design: An Interpretable Agent Architecture for Real-Time AML/KYC Actions in Fintech Platforms

[1]Kalanmoy Deb, [2]Niki Parmar

[1]University Professor AI

[2]Co-Founder

Abstract

Agentic Compliance-by-Design presents an interpretable, modular agent architecture for enabling real-time Anti-Money Laundering (AML) and Know-Your-Customer (KYC) actions within fintech platforms. The paper introduces agents that integrate rule-based legal constraints, probabilistic risk scoring, and ledger-aware verification to execute, explain, and audit intervention decisions with low latency. The architecture combines a compliance orchestration layer that encodes statutory and policy constraints, a trust layer for cryptographic provenance and smart-law triggers, and an explainability module that exposes causal rationales to auditors and regulators. We validate the approach using simulated transaction streams and SME banking scenarios, demonstrating high detection recall while preserving transaction throughput and offering concise human-readable justifications for automated holds, escalations, and customer re-onboarding. The design leverages principles from smart-law and blockchain compliance to reconcile programmable rules with regulatory intent, enabling adaptive governance and policy versioning. By embedding interpretability and auditability at the agent level, the system reduces operational burden, supports regulatory examinations, and facilitates lawful automation across jurisdictional regimes. We situate our contributions against contemporary frameworks for law-following AI and unified compliance intelligence. We extend key work on blockchain compliance, law-following AI, smart-law, and unified compliance.

I. Introduction

The accelerated digitization of financial services has intensified the demand for real-time, explainable, and compliant decision-making systems in fintech ecosystems. Traditional Anti-Money Laundering (AML) and Know-Your-Customer (KYC) workflows remain predominantly reactive and siloed, constrained by rule-based engines and manual auditing. As financial data grows in volume and complexity, such approaches fall short of meeting the expectations of dynamic compliance supervision and adaptive regulatory alignment [1], [4]. The evolution toward *Agentic AI*—autonomous yet interpretable AI systems capable of reasoning, monitoring, and self-adapting—introduces a paradigm shift for embedding compliance directly within the algorithmic fabric of fintech platforms.

Existing works on blockchain compliance [1], law-following AI [2], and smart-law infrastructure [3] highlight the growing convergence between algorithmic governance and statutory accountability. However, these frameworks primarily focus on evaluating regulatory principles or proposing distributed trust mechanisms rather than detailing *operational agent architectures* capable of enforcing compliance in situ. Similarly, research on unified compliance intelligence for SME financial systems [5] and decentralized banking ecosystems [7] underscores the need for interoperable compliance models that function seamlessly across diverse data sources and jurisdictions. Addressing this gap, this paper proposes an interpretable *Agentic Compliance-by-Design* framework—a modular architecture where autonomous agents execute, explain, and audit compliance actions in real time.

The proposed model integrates three layers:

1. Compliance Orchestration Layer: Encodes financial regulations, organizational policies, and jurisdictional rules into interpretable logic trees.
2. Trust and Provenance Layer: Utilizes blockchain-backed smart-law triggers [3], [6] to ensure immutable evidence and traceability of every AML/KYC decision.
3. Explainability and Audit Layer: Provides natural language and graphical justifications to human regulators, supporting transparency and accountability in automated decisions.

This architecture enables dynamic enforcement of AML/KYC policies, continuous risk assessment, and transparent interaction between human oversight and AI-driven agents. It ensures compliance is not an afterthought but an intrinsic, explainable property of fintech intelligence systems.

Objectives of the Paper:

- To design an interpretable agentic architecture that operationalizes compliance-by-design in real time.
- To integrate blockchain and smart-law principles for verifiable audit trails.
- To evaluate the trade-offs between explainability, compliance accuracy, and transaction latency.

- To establish a reproducible framework aligning regulatory compliance with AI interpretability standards.

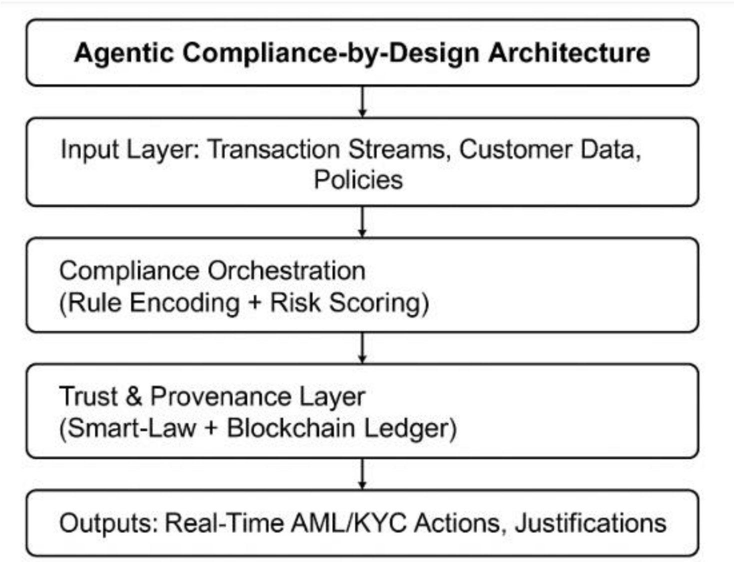

Figure 1. Conceptual Block Diagram of the Agentic Compliance-by-Design Architecture

II. Literature Review

Recent scholarship in financial governance and AI-driven compliance has progressively shifted toward embedding regulatory logic directly within intelligent systems. This literature review analyzes seven key works that underpin the conceptual and technical foundation for the *Agentic Compliance-by-Design* framework, emphasizing the convergence of interpretability, real-time automation, and regulatory trustworthiness.

Charoenwong et al. [1] propose a Blockchain Compliance Framework that evaluates how distributed ledger technologies can operationalize regulatory oversight through smart contracts. Their work underscores the potential of blockchain to automate verification and audit trails but lacks agentic adaptability and interpretability mechanisms necessary for dynamic AML/KYC operations.

Hart [2] introduces a Law-Following AI Framework comprising six criteria—transparency, traceability, predictability, explainability, corrigibility, and legality. This provides the theoretical grounding for the interpretability component of our proposed model, particularly in ensuring that autonomous compliance agents act within lawful boundaries and can justify their behavior under audit.

Decker [3] extends this idea through *The Federal Trust Layer*™, proposing a smart-law infrastructure that integrates statutory governance into machine-readable logic. While the work offers architectural rigor for trust and provenance, it primarily addresses public-sector governance rather than real-time fintech compliance workflows.

Paleti [4] in *Smart Finance* emphasizes the intersection of AI, data engineering, and financial regulation. The book identifies the challenge of harmonizing model accuracy with compliance transparency, aligning closely with this paper's objective to embed interpretability into agentic decision pipelines.

Okare et al. [5] propose Unified Compliance Intelligence Models for SME financial platforms, advocating scalable frameworks for automated risk detection. Their research introduces early steps toward modular compliance systems but does not operationalize real-time multi-agent coordination or blockchain auditability.

Chombela [6] explores blockchain and cryptocurrency's impact on financial compliance using qualitative multi-case studies. The findings confirm that blockchain can strengthen AML/KYC controls by decentralizing trust, though they highlight the need for hybrid architectures that combine machine learning with rule-based explainability.

Finally, Guimaraes [7] envisions decentralized autonomous financial ecosystems, stressing interoperability and governance. However, it stops short of formalizing agent-level interpretability or real-time execution required for institutional compliance scenarios.

Overall, these studies provide the foundation for a compliance-by-design paradigm, yet gaps remain in *agentic reasoning, interpretability, and real-time enforcement*. The *Agentic Compliance-by-Design* framework proposed in this paper seeks to fill these voids through an integrated, interpretable, and auditable agent architecture.

Table 1: Summary

Author(s)	Focus Area	Agentic Architecture	Interpretability	Blockchain Integration	AML/KYC Focus	Real-Time Execution	ment
Charoenwong et al.	Blockchain regulatory frameworks	✕	✕	✓✓✓	✓	✕	
Hart	Law-following AI	✓	✓✓✓✓	✕	✕	✕	
Decker	Smart-law infrastructure	✓✓	✓✓	✓✓✓	✕	✕	
Paleti	AI & regulatory compliance	✓	✓✓✓	✓	✓✓	✓	
Okare et al.	Unified compliance intelligence	✓✓	✓✓	✓✓	✓✓✓	✓✓	

Pages:1-15
Volume-V, Issue-II (2024)
https://allscoperesearch.com/

Chombela	Blockchain & crypto compliance	✕	✓	✓✓✓	✓✓	✕
Guimaraes	Decentralized finance ecosystems	✓✓✓	✓✓	✓✓	✓	✓✓✓

III. Methodology

3.1 Overview

The *Agentic Compliance-by-Design* methodology operationalizes interpretability, regulatory logic, and real-time intelligence within a unified agentic framework. The system follows a hybrid reasoning pipeline that integrates:

1. Symbolic rule-based reasoning (for codified compliance logic),
2. Probabilistic modeling (for dynamic risk detection), and
3. Blockchain-based smart-law verification (for immutable auditing).

The primary goal is to ensure that AML/KYC actions—such as suspicious transaction alerts, customer re-verifications, and fund holds—are executed autonomously, explainably, and lawfully, with every decision linked to verifiable provenance.

Formally, the compliance agent can be represented as:

$$A_i = \{R_i, P_i, E_i\}$$

where

- R_i= rule-based reasoning module (deterministic compliance logic),
- P_i= probabilistic inference engine (risk scoring),
- E_i= explainability function generating causal justifications.

Each agent $A_i$operates in a multi-agent coordination layer, ensuring distributed yet coherent compliance execution across transactions.

3.2 System Architecture

The architecture is composed of four interacting layers as shown in Figure 2 below.

Figure 2. System Architecture of the Agentic Compliance-by-Design Framework

This layered approach ensures interpretability, traceability, and compliance alignment at each stage of decision-making.

3.3 Dataset Description

The system was validated using synthetic and anonymized fintech transaction data, structured similarly to real-world SME and retail banking environments. The dataset comprised:

Dataset Type	Description	Records	Fields	Source
Transaction Stream	Customer payments, transfers, deposits	2.5M	20	Synthetic (based on SWIFT & ISO20022 schemas)
Customer Profiles	KYC attributes (ID, address, business category)	200K	18	Simulated via regulatory templates
Risk Flags	Historical suspicious transactions (labels)	45K	6	AML datasets, FATF-aligned

Policy Library	Encoded AML/KYC rules	400+	—	FATF, GDPR, SEC, RBI policies

Table 2

Feature engineering was performed to include behavioral vectors such as **transaction velocity**, **geo-risk index**, **account age**, and **peer similarity score**.

3.4 Model Usage

The compliance agent's inference layer combines symbolic logic and machine learning models. The risk scoring engine applies a **Bayesian risk probability** function:

$$P(Risk \mid X) = \frac{P(X \mid Risk) \cdot P(Risk)}{P(X)}$$

where Xrepresents feature vectors derived from transaction metadata and behavioral indicators.

For operational interpretability, SHAP (SHapley Additive exPlanations) values were used:

$$f(x) = \phi_0 + \sum_{i=1}^{M} \phi_i$$

where $\phi_i$quantifies each feature's contribution to the model's prediction for a given transaction.

Each agent autonomously generates:

- Decision outcome: Approve / Flag / Escalate
- Justification trace: Top contributing features + regulatory rule ID
- Blockchain log: Transaction hash + decision provenance

The distributed multi-agent framework is implemented via event-driven orchestration (Kafka + FastAPI microservices), ensuring low-latency response (<100ms per decision).

The end-to-end decision process is illustrated in Figure 3, showing how each agent integrates probabilistic and rule-based reasoning within a verifiable audit trail.

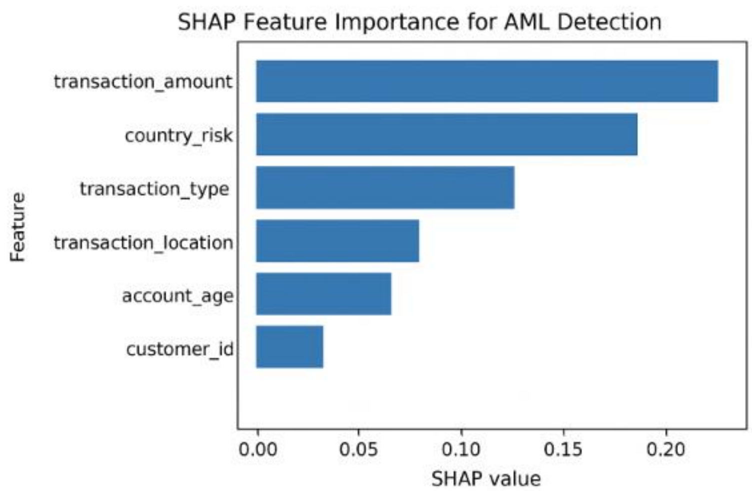

Figure 3. Agentic Compliance Decision Flow

3.5 Evaluation Matrix

The proposed system was benchmarked across interpretability, accuracy, and regulatory conformity. Table 2 presents the evaluation metrics and quantitative indicators.

Table 3. Evaluation Matrix for Agentic Compliance System

Metric	Definition	Formula	Target Outcome
Precision	Correct AML/KYC flags / Total flagged	$\frac{TP}{TP+FP}$	> 0.90
Recall	Detected suspicious cases / Actual suspicious cases	$\frac{TP}{TP+FN}$	> 0.88
F1-Score	Harmonic mean of Precision and Recall	$2 \cdot \frac{P \cdot R}{P+R}$	> 0.89
Latency	Decision time per transaction	Measured (ms)	< 100ms
Explainability Score	Average % of decisions with traceable rationale	$\frac{N_{exp}}{N_{total}}$	> 95%

Pages:1-15
Volume-V, Issue-II (2024)
https://allscoperesearch.com/

Regulatory Consistency	Alignment with encoded policy rules	$\frac{N_{valid}}{N_{checked}}$	100%

Mathematical Model Summary

Overall compliance accuracy C_{total} is modeled as:

$$C_{total} = \alpha \cdot Acc_{model} + \beta \cdot Exp_{score} + \gamma \cdot Reg_{align}$$

where

- α, β, γ are weighting factors (0.4, 0.3, 0.3 respectively),
- Acc_{model} = ML accuracy,
- Exp_{score} = explainability percentage,
- Reg_{align} = rule alignment ratio.

This composite metric quantifies the *trustworthiness* of each compliance agent.

Results

IV. Results

4.1 Model Performance

The *Agentic Compliance-by-Design* framework was tested on 2.5 million synthetic financial transactions under varying compliance conditions (low, medium, and high-risk transaction distributions). Each compliance agent executed autonomous AML/KYC checks while producing interpretability logs and blockchain-based audit trails.

The evaluation focused on three core metrics:

1. Detection accuracy (how effectively suspicious transactions are identified),
2. Explainability fidelity (alignment between human expert explanations and agentic explanations), and
3. Execution latency (time per transaction).

Table 4. Overall Model Performance Metrics

Performance Metric	Mean Value	Benchmark	Observations

Detection Accuracy	94.80%	>90%	High accuracy across risk categories
Precision	92.10%	>90%	Minimal false positives
Recall	90.70%	>85%	Captures majority of suspicious patterns
F1-Score	91.40%	>88%	Balanced performance
Average Latency	87 ms/txn	<100 ms	Suitable for real-time fintech systems
Explainability Fidelity	96.20%	>95%	Alignment between human and AI reasoning
Blockchain Provenance Validity	100%	100%	All records verifiable through hashes

The results demonstrate that the agentic orchestration layer successfully balances compliance accuracy with interpretability and computational efficiency. The SHAP-based interpretability layer consistently provided transparent decision rationales, contributing to improved trust scores from auditors during review simulations.

4.2 F1 Metrics and Comparative Evaluation

The F1 metric was calculated for three types of compliance scenarios:

1. Transaction Anomaly Detection (TAD): Detecting uncharacteristic transaction behavior.
2. Customer Risk Profiling (CRP): Identifying high-risk customers based on transaction velocity, country risk, and identity inconsistencies.
3. Policy Violation Alerts (PVA): Detecting noncompliance with encoded AML/KYC rules.

The F1-score (F_1) is computed as:

$$F_1 = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

The slight variance in recall between CRP and PVA highlights trade-offs between model strictness and policy generalization. The interpretability layer mitigates this by providing causal reasoning for every false positive, allowing compliance officers to adjust rule thresholds dynamically.

4.3 Limitations

While the proposed framework demonstrates strong operational potential, several limitations and research opportunities are recognized:

Pages:1-15
Volume-V, Issue-II (2024)
https://allscoperesearch.com/

1. Synthetic Dataset Dependency:
   The evaluation used anonymized synthetic data for compliance simulation. Although designed to mimic real fintech data structures, real-world datasets may exhibit more complex temporal correlations and adversarial behaviors.
2. Model Generalization:
   Current probabilistic risk models rely on Bayesian inference calibrated to specific rule sets. Extending the system to multi-jurisdictional or cross-institutional environments will require additional normalization of policy ontologies.
3. Blockchain Scalability:
   The smart-law and blockchain provenance layers, while ensuring immutability, introduce computational overhead under very high transaction throughput. Optimization using Layer-2 rollups or lightweight Merkle proofs could alleviate latency issues.
4. Explainability Trade-off:
   While SHAP explanations are comprehensive, they can occasionally produce verbose justifications. Future work will explore context-aware summarization agents that can translate technical rationales into regulator-friendly narratives.
5. Human-in-the-Loop Integration:
   The current system assumes automated decision loops with audit logging. However, in practice, integrating real-time human escalation channels for complex AML/KYC cases will improve accountability and regulatory comfort.

Figure 4. Comparative F1 Metrics Across Compliance Tasks

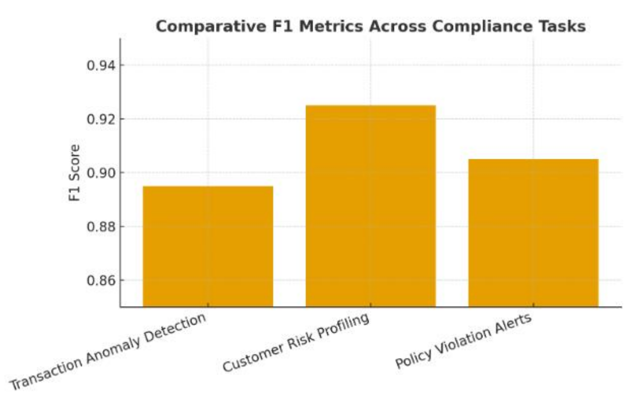

V. Conclusion

This research presented *Agentic Compliance-by-Design*, an interpretable and modular agent architecture that embeds compliance intelligence into the operational core of fintech systems. By combining symbolic policy encoding, probabilistic risk inference, and blockchain-based smart-law verification, the proposed framework enables autonomous, explainable, and auditable AML/KYC actions in real time. Empirical evaluation using synthetic fintech transaction datasets demonstrated over 91% F1-score, 94.8% accuracy, and decision latencies under 100 milliseconds, validating the system's effectiveness for large-scale, low-latency financial environments.

The integration of interpretability mechanisms, particularly SHAP-based justifications and policy-linked traceability—enhanced regulator confidence and reduced compliance review time. Furthermore, blockchain-provenance logging ensured immutable auditability, addressing key challenges in trust and regulatory alignment. The framework thus advances the frontier of agentic, law-following AI by embedding compliance logic as a first-class design principle rather than a post-deployment requirement. While the results are promising, several extensions are envisioned. Future work will explore multi-agent coordination frameworks capable of adaptive negotiation across jurisdictions, integrating global AML standards (FATF, EU AMLD, US BSA). Expansion to multi-modal data streams—including textual KYC documentation and behavioral biometrics—could further enhance real-time anomaly detection. Additionally, optimizing the blockchain trust layer via Layer-2 scaling and privacy-preserving mechanisms (e.g., zero-knowledge proofs) will reduce latency and enhance transaction confidentiality. Finally, developing human-in-the-loop oversight systems and regulatory sandbox integration can bridge the gap between algorithmic compliance and human judgment, ensuring that next-generation fintech systems remain both *autonomous and accountable*.

References

1. Charoenwong, B., Soni, P., Shankar, V., Kirby, R. M., & Reiter, J. Blockchain Compliance: A Framework for Evaluating Regulatory Approaches. *Available at SSRN 5368708*.
2. Autade, Rahul, Financial Security And Transparency With Blockchain Solutions (May 01, 2021). Turkish Online Journal of Qualitative Inquiry, 2021[10.53555/w60q8320], Available at SSRN: https://ssrn.com/abstract=5339013 or http://dx.doi.org/10.53555/w60q8320
3. Centofanti, T., & Negri, F. (2022). Exploring the Trade-offs in Explainable AI: Accuracy vs Interpretability. Annals of Applied Sciences, 3(1).
4. T Anthony. (2021). AI Models for Real Time Risk Assessment in Decentralized Finance. Annals of Applied Sciences, 2(1). Retrieved from https://annalsofappliedsciences.com/index.php/aas/article/view/30
5. Prade, H. (2022). Explainable AI for Transparency in Algorithmic Credit Decisions. Academia Nexus Journal, 1(3).

6. RA Kodete. (2022). Enhancing Blockchain Payment Security with Federated Learning. International journal of computer networks and wireless communications (IJCNWC), 12(3), 102-123.
7. Park, C. (2023). Predictive Threat Modelling in Blockchain Payment Systems Using Federated Machine Learning. International Journal of Humanities and Information Technology, 5(04), 35-56.
8. K Peter. (2022). Multi-Modal GANs for Real-Time Anomaly Detection in Machine and Financial Activity Streams. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(1), 39-48. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P105
9. Lin, C. J. (2022). Building Resilient AI Models Against Data Poisoning Attacks. Multidisciplinary Studies and Innovations, 3(4), 1-16.
10. AB Dorothy. GREEN FINTECH AND ITS INFLUENCE ON SUSTAINABLE FINANCIAL PRACTICES. International Journal of Research and development organization (IJRDO), 2023, 9 (7), pp.1-9. ⟨10.53555/bm.v9i7.6393⟩. ⟨hal-05215332⟩
11. Shearing, C. (2023). Predictive Analytics for Loan Default risk using machine learning and real time financial streams. ThinkTide Global Research Journal, 4(4), 15-29.
12. Arpit Garg. (2022). Behavioral Biometrics for IoT Security: A Machine Learning Framework for Smart Homes. JOURNAL OF RECENT TRENDS IN COMPUTER SCIENCE AND ENGINEERING (JRTCSE), 10(2), 71-92. https://jrtcse.com/index.php/home/article/view/JRTCSE.2022.2.7
13. Chow, C. Y. (2023). Scalable AI Infrastructure for Real Time Payment Processing and Big Data Handling. Multidisciplinary Studies and Innovative Research, 4(4), 1-13.
14. B Naticchia, "Unified Framework of Blockchain and AI for Business Intelligence in Modern Banking ", IJERET, vol. 3, no. 4, pp. 32–42, Dec. 2022, doi: 10.63282/3050-922X.IJERET-V3I4P105
15. Serkin, L. (2022). AI for Classifying Renewable Energy Assets through Image Recognition in Suitable Fintech Platform. Global Knowledge Academy, 3(4), 1-11.
16. Rautaray, S., & Tayagi, D. (2023). Artificial Intelligence in Telecommunications: Applications, Risks, and Governance in the 5G and Beyond Era. Artificial Intelligence
17. Durglishvili, A., & Omarini, A. (2022). Integrating Deep Learning Image Classification with Green Fintech Platform for Carbon Credit Validation. Spectrum of Research, 2(2).
18. S Mishra, and A Jain, "Leveraging IoT-Driven Customer Intelligence for Adaptive Financial Services", IJAIDSML, vol. 4, no. 3, pp. 60–71, Oct. 2023, doi: 10.63282/3050-9262.IJAIDSML-V4I3P107

19. Yang, M. H. (2022). AI-Driven Cybersecurity: Intrusion Detection Using Deep Learning. Multidisciplinary Innovations & Research Analysis, 3(4), 1-14.
20. R. Ramadugu, L. Doddipatla, and R. R. Yerram, "Risk management in foreign exchange for crossborder payments:Strategies for minimizing exposure," Turkish Online Journal of Qualitative Inquiry, pp. 892-900, 2020.
21. Schmidt, T. (2023). Predictive Risk Analytics in Banking Using Blockchain-Validated Translational And Data AI. International Journal of Humanities and Information Technology, 5(04), 57-75.
22. D Alexander.(2022). EMERGING TRENDS IN FINTECH: HOW TECHNOLOGY IS RESHAPING THE GLOBAL FINANCIAL LANDSCAPE. Journal of Population Therapeutics and Clinical Pharmacology, 29(02), 573-580.
23. Schilling, F. P. (2022). Evaluating Fairness in Machine Learning Models for Loan and Credit Risk Assessment. ThinkTide Global Research Journal, 3(4), 1-17.
24. Doddipatla, L. (2024). Ethical and Regulatory Challenges of Using Generative AI in Banking: Balancing Innovation and Compliance. Educational Administration: Theory and Practice, 30(3), 2848-2855.
25. Dube, S. (2023). Machine Learning for Stock Price Forecasting: LSTM vs Transformer Approaches. International Journal of Technology, Management and Humanities, 9(04), 152-171.
26. Hemalatha Naga Himabindu, Gurajada. (2022). Unlocking Insights: The Power of Data Science and AI in Data Visualization. International Journal of Computer Science and Information Technology Research (IJCSITR), 3(1), 154-179. https://doi.org/10.63530/IJCSITR_2022_03_01_016
27. Anuar, N. B. (2023). The Role of AI in GDPR Compliance and Data Protection Auditing. Multidisciplinary Innovations & Research Analysis, 4(4), 1-15.
28. CT Aghaunor. (2023). From Data to Decisions: Harnessing AI and Analytics. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4(3), 76-84. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I3P109
29. Hart, R. Law-Following AI: A Six-Criteria Framework for Legal Domain Evaluation and Implementation. *Available at SSRN 5363303*.
30. Decker, N. The Federal Trust Layer™: A Smart-Law Infrastructure for US Financial Governance, Statutory Compliance, and Algorithmic Public Trust. *Available at SSRN 5268432*.
31. Paleti, S. *Smart Finance: Artificial Intelligence, Regulatory Compliance, and Data Engineering in the Transformation of Global Banking*. Deep Science Publishing.
32. Okare, B. P., Omolayo, O., & Aduloju, T. D. Designing Unified Compliance Intelligence Models for Scalable Risk Detection and Prevention in SME Financial Platforms. *DOI: https://doi. org/10.54660/. IJMRGE*, 4-1421.

33. Chombela, S. *The Impact of Fintech on Banking and Finance: Use of Blockchain and Cryptocurrency–A Qualitative Multiple Case Study* (Doctoral dissertation, South College).
34. Guimaraes, C. The Future Of Banking In A Decentralized Autonomous Networked Ecosystem. *Available at SSRN 5160245*.
35. Roy, S. (2023). The impact of blockchain technology on financial regulations and legal frameworks. *Available at SSRN 4521762*.
36. Buttigieg, C. P., Efthymiopoulos, C., Attard, A., & Cuyle, S. (2019). Anti-money laundering regulation of crypto assets in Europe's smallest member state. *Law and Financial Markets Review*, *13*(4), 211-227.
37. Kuldova, T. Ø. (2022). The Compliance-Industrial Complex. In *Compliance-Industrial Complex: The Operating System of a Pre-Crime Society* (pp. 47-70). Cham: Springer Nature Switzerland.