Amit Butail, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science,
Engineering and Technology
An Open Access Journal

AI-Driven Credit Scoring in Hybrid Cloud Banking: An Investigative Assessment of Explainability and Regulatory Transparency Gaps

Amit Butail, Anamika
Yogananda School of AI, Computers and Data Sciences, Shoolini University, Solan 173229, India,

Abstract- The introduction of Artificial Intelligence (AI) into the credit scoring process has transformed banks' capacity to evaluate risk, however, the deployment of AI models in a hybrid cloud environment brings significant challenges of transparency and explainability in terms of regulations. This study examines the current state of AI credit scoring systems in hybrid cloud banking systems and systematically pinpoints the regulatory transparency challenges that hinder compliance with the banking regulations in the United States, including the Federal Reserve's Model Risk Management guidance (SR 11-7), the Federal Housing Administration's Fair Housing Act (FHA), the Equal Credit Opportunity Act (ECOA), and the National Institute of Standards and Technology's AI Risk Management Framework. To investigate the operational architecture, explainability challenges, and compliance issues of hybrid cloud AI credit scoring systems, a systematic literature review was performed based on the PRISMA framework, which included peer-reviewed publications, regulatory documents, and industry reports from 2020 to 2025. Here are three areas where we found insufficient transparency: (i) the "black box" issue with tracing back opaque decision-making processes due to Federal Reserve Board SR 11-7 remains unaddressed, (ii) lack of explainability mechanisms originating from consumer protection frameworks under the FHA and ECOA, and (iii) lack of transparency and accountability frameworks with AI models in distributed hybrid cloud systems. Therefore, the research contributes to a comprehensive classification of transparency and explainability gaps in hybrid cloud credit scoring, particularly in relation to US regulatory frameworks, and provides basic building blocks for explainable AI (XAI) systems that comply with regulatory requirements.

Keywords: Artificial Intelligence, Credit Scoring, Hybrid Cloud Banking, Explainable AI, Regulatory Compliance, Model Risk Management, Fair Lending, SR 11-7.

I. INTRODUCTION

The integration of AI and hybrid cloud technologies has the potential to reshape credit scoring by processing large volumes of structured and unstructured data at a greater level of predictive accuracy [1], [2]. Financial services organizations are beginning to recognize this opportunity. When hybrid cloud environments are utilized in combination with AI-backed scoring models, banks achieve greater flexibility, quicker iterative cycles for model development, and distributed analytics [3], [4]. Simultaneously, the adoption of this technology creates certain challenges for regulatory disclosures. Some examples of these challenges are the ambiguous nature of model usage, the fragmented nature of data processing, and the complicated nature of auditing [5].

Traditional credit scoring approaches employing the FICO scoring model, for instance, have specified rules and are more bureau-centric, less flexible credit variables. However, contemporary AI systems leverage machine learning algorithms like ensemble learners, neural networks, and transformer-based large language models (LLMs) to mine information from multiple behavioral and transactional data sources as well as alternative data sources [6],[7].

Existing research shows that AI-based risk modeling can significantly outperform linear approaches, but adds opacity that has made it difficult to meet supervisory requirements, adverse action disclosure and fairness requirements [8], [9], [10]. While some explainable AI (XAI) techniques, like LIME, SHAP, surrogate models or counterfactual reasoning, offer some interpretability, it has been demonstrated that most of these techniques produce technical feature importance explanations which do not satisfy customers, auditors, and regulators [12], [13].Recent research suggests structured XAI frameworks which explicitly map the behavior of the model to a human-interpretable path of reasoning that can provide a more transparent foundation for communications with auditors, regulators, and customers [14].

The move to hybrid cloud adds to the transparency issues. Banks store personal and financial information in private or on-premises systems, and use the public cloud for computationally demanding tasks such as model training, large-scale model inference and enforcing rules in real-time [15]. This architecture is resilient and scalable, but it also creates a fragmented decision pipeline, log repositories, and decision governance controls, and it can be challenging to show end-to-end traceability or trace decision logic back to the source where it can be reproduced [3]. Existing regulations, such as those from the Bank for International Settlements (BIS), stress the need for these distributed model environments to be more auditable, explainable, and accountable [4].

Although there is as yet little systematic research on the particular regulatory transparency challenges that emerge when AI credit scoring systems and hybrid cloud architectures intersect, the technical capabilities of AI systems and the operational advantages of hybrid cloud architectures have been well documented. Previous works tend to focus on explainability issues without considering deployment architectures, while other works focus on governance issues for hybrid cloud architectures without considering the requirements for transparency of AI.

This research addresses this gap which aims to explore: (i) The current state of the credit scoring mechanisms for hybrid cloud banking systems, their architecture patterns, data flow designs, and operational workflows. (iii) Identifying and mapping regulatory transparency gaps in current implementations to US banking regulations (SR 11-7, FHA, ECOA, NIST AI RMF). (iii) The explainability gaps for from meeting regulatory consumer protection requirements in hybrid cloud AI credit scoring systems.

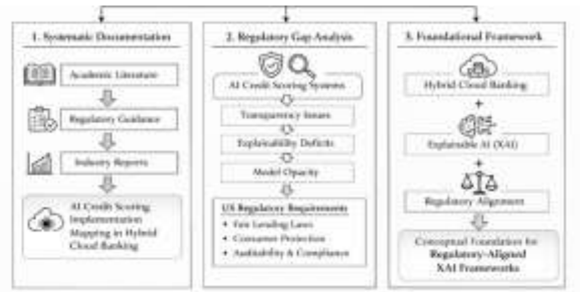

Figure 1: Study contributions illustrating systematic documentation, regulatory gap analysis, and the development of an explainable AI framework for hybrid cloud banking credit scoring.

The three main contributions of the study are shown in Figure 1. It first identifies the mapping of AI-based credit scoring by conducting a literature review, regulatory guidelines and industry reports. It then highlights regulatory gaps with regard to regulatory requirements, including transparency, lack of explainability and model opacity. Lastly, it suggests a basic Explainable AI (XAI) framework to enhance the alignment of regulations and build trust in the use of AI for credit decisions in the context of hybrid cloud banking. The results will be the basis for the XAI framework that aligns with the regulatory landscape and is adapted to hybrid cloud financial ecosystems.

II. RESEARCH METHODOLOGY

In this study, a systematic literature review methodology based on the PRISMA protocol is used along with design science preparatory analysis to thoroughly explore the implementation of AI-based credit scoring in hybrid cloud banking environments and to uncover the regulatory transparency gaps. The design ensures that the research is comprehensive and that the synthesis and analysis are conducted in a research-oriented manner.

Research Design

The research design contains the following four interrelated components:

1. The systematic literature review (SLR) that is anchored in the PRISMA procedure for the identification, review, and synthesis of peer-reviewed scholarly articles, regulatory literature, and catalogs of industry literature;

2. The regulatory mapping framework that comprises the systematic identification and coding of relevant provisions of US banking regulations (i.e. SR 11-7, FHA, ECOA, NIST AI RMF) and their structured compliance articulations;
3. The thematic synthesis of transparency and explainability gaps that involves layered coding and classification; and
4. Deliberate and purposeful gap analysis of the compliance mapping of observed implementation methodologies and the required regulatory compliance frameworks.

Search Strategy

Figure 2 demonstrates the systematic approach to the structured search used when compiling studies from academic databases, regulatory repositories and industry sources. Integrating these sources ensures the most complete and precise literature on AI credit scoring and hybrid cloud banking. The pathways and protocols used in the databases are displayed in the search strategy diagram, while the search outlines the three primary sources, academic, regulatory and industry, to assess the literature and conduct a robust and even descriptive study.

Among academic sources databases such as IEEE Xplore, ACM Digital Library, Scopus, Web of Science, and Google Scholar. These databases provides credit scoring, cloud computing, and AI peer-reviewed articles and research studies. Regulatory repositories comprise publications from the Federal Reserve Board, the Office of the Comptroller of the Currency (OCC), the Consumer Financial Protection Bureau (CFPB), and NIST.

These sources offer official guidelines and regulatory requirements regarding the financial system, risk management, and AI transparency. The industry sources consist of technical documentation from major cloud providers (AWS, Google Cloud, Microsoft Azure), financial technology white papers, and banking industry reports. These sources offer practical information on actual application of AI-powered credit-scoring systems.

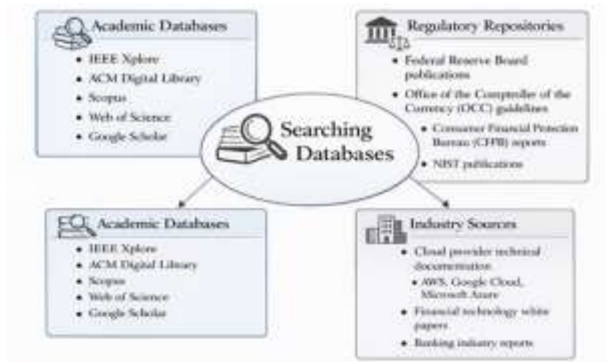

Figure 2: Literature search strategy illustrating academic, regulatory, and industry databases used for systematic evidence collection.

Primary search strings combined Boolean operators and domain-specific terminology:
(artificial intelligence OR machine learning OR deep learning OR neural networks OR LLM OR large language model) AND (credit scoring OR credit risk OR lending OR creditworthiness) AND (hybrid cloud OR cloud computing OR cloud deployment OR multi-cloud) ("banking" OR "financial services" OR "fintech") AND ("explainability" OR "interpretability" OR "transparency" OR "XAI" OR "explainable AI" OR "model risk management" OR "SR 11-7" OR "fair lending"). The literature was compiled from January 2020 to December 2025, with a specific focus on publications from 2024-2025, which correspond to the latest regulatory developments and technological advances in credit scoring with the use of LLMs.

Inclusion and Exclusion Criteria

The studies selected were screened for relevance, quality, and to meet the goals of this systematic review. To ensure that only studies related to credit scoring systems using AI in hybrid cloud banking environments are included, the inclusion and exclusion criteria were defined as shown in Table 1.

Table 1: Inclusion and Exclusion Criteria for Study Selection.

Inclusion Criteria:	Exclusion Criteria:
Peer-reviewed journal articles, conference proceedings, and regulatory publications addressing AI/ML in credit scoring	Publications focused exclusively on non-banking financial services (insurance, securities trading)

	without credit scoring relevance
Studies examining hybrid cloud, multi-cloud, or cloud-native banking architectures	Studies addressing only on-premises AI systems without cloud deployment considerations
Publications discussing explainability, interpretability, or transparency in financial AI systems	Publications lacking data, architectural descriptions, or regulatory analysis
Regulatory guidance documents specific to US banking (Federal Reserve, OCC, CFPB, NIST)	Opinion pieces, blog posts, and non-peer-reviewed sources lacking technical rigor
Industry reports from established cloud providers and financial institutions documenting implementation practices	Duplicate publications or conference papers superseded by journal versions
Studies published in English language	Publications focused exclusively on non-banking financial services (insurance, securities trading) without credit scoring relevance

Study Selection Process

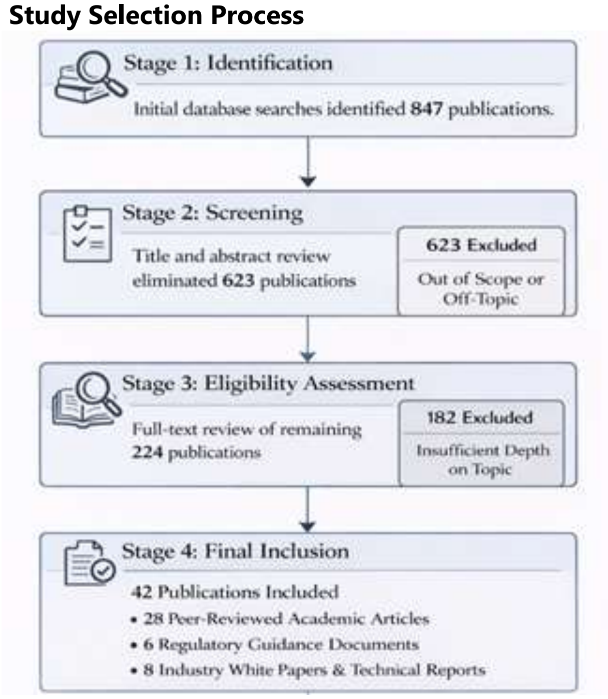

Figure 3: PRISMA flow diagram illustrating the systematic study selection process from initial identification to final inclusion of relevant publications.

Figure 3 uses a PRISMA-compliant flow diagram for the systematic review and shows how studies were selected. During the identification stage, a total of 847 records were found, of which 623 were removed during title and abstract screening. 224 studies were included for full-text assessment. 182 studies were removed after this assessment, and 42 studies were selected. To ensure a descriptive balance, these studies comprise industry documents, regulatory reports, and academic studies.

Data Extraction and Synthesis

The studies we chose helped us pull out details about implementation, characteristics, explainability techniques, regulations, and research gaps. What we collected was analyzed and compared to determine if there were common themes and emerging AI technologies and transparency challenges related to credit scoring within hybrid cloud banking.

Table 2 shows the key dimensions that were extracted and synthesized from selected studies, such as: implementation characteristics, explainability mechanisms, regulatory dimensions, and identified research gaps and future directions.

Table 2: Framework for data extraction and synthesis across implementation, explainability, regulatory, and research gap dimensions.

Implementation Characteristics:	Explainability Mechanisms:	Regulatory Dimensions:	Research Gaps and Future Directions:
AI/ML model types and architectures	XAI techniques employed (LIME, SHAP, attention mechanisms, etc.)	Specific compliance requirements addressed	Explicitly stated limitations
Hybrid cloud deployment patterns	Transparency features and audit trail capabilities	Documentation and validation practices	Recommended research priorities

Data pipeline designs and security controls	Stakeholder-facing explanation interfaces	Governance frameworks and oversight mechanisms	Emerging technology considerations (LLMs, RAG, agentic AI)
Integration approaches with banking systems		Identified compliance gaps or challenges	

Thematic Analysis Framework

A thematic coding method was used in an iterative process to identify patterns, gaps in the literature and priorities in the selected literature. The thematic areas of interest for the analysis were the evolution of AI applications in banking compliance, technical architectures in hybrid cloud credit scoring system, challenges of explainability and interpretability of AI models, requirements of the regulatory framework and their evolution, and identification of research gaps in transparency mechanisms.

The coding process was systematic and included several stages that were structured. Open coding was first done and 87 concepts were identified from the literature reviewed. These concepts were then grouped into meaningful sub-themes using axial coding (23 sub-themes, based on conceptual relationships). These sub-themes were then used in selective coding to further narrow them down to five major themes that were the main analytical dimensions of the study. To categorize these themes, an independent review was conducted to increase the study's reliability and reduce bias.

Regulatory Mapping Methodology

To analyze US banking regulations, we mapped the operational and technical requirements for the following:

- **SR 11-7 Model Risk Management:** Documentation, Validation, and Monitoring Standards and Model Catalog Requirements
- **FHA/ECOA Fair Lending:** Anti-Discrimination, Adverse Action Notifications, and Disparate Impact
- **NIST AI RMF:** Transparency, Trust, and Risk

These were compared against the AI credit scoring implementation patterns, and uncovered some specific compliance concerns.

Limitations and Mitigation

The most significant limitations for this study are the rapid pace of the development of AI technologies and time-sensitive findings. Moreover, our lack of access to proprietary banking systems and internal credit scoring systems constrains us from anything beyond publicly available research. There is consistent shift in the AI regulation in financial services, and future iterations of regulation may provide regulatory requirements beyond the current study.

To address these challenges, this research targets the most recent publications from 2024-2025 to help account for the newest technological and regulatory developments. To assist in ensuring the accuracy of the findings, these were triangulated in the context of the literature, regulatory standards, and industry reports. Furthermore, the findings consider the rapidly changing regulatory environment to incorporate flexibility in the findings and the timeframe in which they may change.

3. AI Driven Credit Scoring in Hybrid Cloud Banking: Current State Analysis

This part details the results of the systematic literature review focusing on the contemporary adoption of AI credit scoring within the hybrid cloud banking domain, including the ways in which design, operation, and deployment may be expressed.

Evolution of AI Credit Scoring Implementations

The adoption of AI within banking credit scoring has been segmented into four phases, each characterized by a specific technological paradigm, principal use cases, and associated challenges, as summarized in Table 3 below.

Table 3: Evolution of AI In Banking Credit Scoring

Stage	Dominant AI Methods	Primary Credit Applications	Limitations

Detection & Monitoring (2015-2018) [4]	Rule-based ML, Decision Trees	Credit risk flagging, fraud rules	High false positives, static logic, limited feature complexity
RegTech Automation (2018-2021) [17]	Supervised ML, early NLP	Risk scoring automation, document processing, compliance monitoring	Limited interpretability, siloed systems, regulatory skepticism
Hybrid Cloud-Native (2021-2023) [3], [5]	Deep learning, transformers	Real-time credit decisioning, continuous monitoring, anomaly detection	Opaque decisions, cloud governance issues, audit trail gaps
Generative & Agentic (2023-Present) [18], [19]	LLMs, GenAI, RAG architectures	Automated credit assessment, unstructured data analysis, adaptive risk modeling	Explainability deficits, provenance complexity, regulatory uncertainty

The four phases of AI technology adoption within banking credit scoring are summarized in Table 3. The phases capture the movement from rudimentary techniques like rules-based systems and decision trees, to complex technologies such as deep learning, transformers, and generative AI. The use cases and the technology for the scoring systems have advanced from detecting fraudulent transactions to assessing credit and making instant decisions. With this advancement, new challenges such as interpretability, governance, and regulatory compliance have emerged.

Technical Architecture Patterns

The current hybrid cloud banking implementations of AI credit scoring has three predominant architectural styles, each offering a different level of transparency:

Distributed Inference Architecture

This method separates the training of the model (done on-premise or in a private cloud) from the serving of inference (done in the public cloud). In this case, the credit scoring models are trained in banks' secure environments, and are then exported and implemented as containerized microservices in the public cloud, which provide the necessary scalability for inference [15]. This architecture ensures data sovereignty, reduces the cost of computation, and offers elastic scaling for large scale operations. Also, it creates transparency issues like inability to trace model provenance across environments, inconsistent audit logging, and heterogeneous governance models [16].

Federated Learning Architecture

Models can be trained collaboratively in the presence of data silos without the need to share sensitive credit data. In this case, each bank branch or business unit trains a local model, and encrypted model gradients are sent to a central unit, which builds the consensus model [20]. This not only ensures data privacy but also facilitates regulatory adherence by minimizing data, and allows for collaboration among various institutions. It poses transparency problems like the inability to trace decisions to a particular source of training data, to understand aggregated model behavior, and to keep track of local training processes.

Hybrid RAG Architecture

Today, RAG systems are often adopted to enrich credit scoring models with information from enterprise knowledge bases, such as regulatory documents, credit policies, and past lending decisions, among others [21]. Theoretically, LLM-generated credit assessments and grounded outputs in retrieved organizational knowledge will enhance transparency, with the potential of citation mechanisms. This architecture provides a semantic interpretation of unstructured credit applications, facilitates policy aligned decision reasoning, and dynamically adjusts to regulatory changes. It brings in transparency problems like hallucination in the outputs of the LLMs, the ranking mechanism is opaque, the generated explanations are hard to validate and the provenance is complex and multi-document based [22].

AI Model Types in Production Credit Scoring

The AI model types that have been found to be most widely used in production credit scoring systems are as follows:

Ensemble Methods (Gradient Boosting, Random Forests)

These are the most widely used AI models in production credit scoring systems, as they are generally more effective when applied to tabular credit data and are also more interpretable than neural networks [24], [20]. However, even ensemble models come with their own set of problems in terms of explaining the model at the feature level, which is understandable by non-technical stakeholders.

Deep Neural Networks

They are used more and more in high dimensional analysis of alternative credit data, but their opacity makes them hard to be accepted for regulatory purposes [20]. It has been found that "interpretability concerns related to ANN components need to be addressed to ensure regulatory transparency" [20].

Hybrid SEM-ANN Architectures

Newer techniques include a hybrid of SEM and ANN, which have been used to achieve a compromise between predictive ability and causal interpretation, but have not been widely adopted [20].

Transformer-Based LLMs

Recent implementations have used FinBERT, FinGPT, and domain-adapted language models to obtain credit-relevant information from unstructured application documents, but "the interpretability of these is still limited" [7]. The need for interpretability mechanisms, standardized evaluation and robustness checks for LLM-based credit risk systems is highlighted in a systematic review [7].

Explainability Mechanisms in Current Implementations

There are limited and inconsistent explainability mechanisms in current implementations despite the requirement for transparency in the regulations:

Post-Hoc Explainability Techniques:

- **SHAP (SHapley Additive exPlanations):** Most widely implemented technique for obtaining feature importance scores for single predictions [25], [11]. Real-time implementation of SHAP explanations remains a challenge due to the computational burden and stability issues [13]. Cost-sensitive ML investigations show that while these explanations help in the identification of important predictors, the existing implementations of SHAP explanations do not meet the standards of regulatory interpretability or the reasoning required to comply with the provision of adverse action notifications [26].
- **LIME (Local Interpretable Model-agnostic Explanations):** used for the generation of local approximations of complex models, albeit with questionable fidelity to the actual model [11].
- **Attention Mechanisms:** Transformer models are expected to associate attention weights to features. In reality, attention weights do not really reflect the reasoning behind the model's decision [7].

Inherent Interpretability Approaches:

- **Linear Models with Regularization:** Some organizations have interpretable baseline models, alongside complex AI models, for regulatory reporting. This, however, is a two-model approach, and therefore, consistency risks are an issue [27].
- **Decision Tree Ensembles with Depth Constraints:** Reduced tree size to make trees easier to interpret, with a tradeoff in predictive accuracy [27].

Current implementations of explainability do not adhere to regulatory requirements for full audit trails. As mentioned in recent research, the current explainability approaches offer rudimentary explanations of model behaviour, but future research should aim to create more fine-grained and accessible explainable AI (XAI) modules for non-technical stakeholders like regulators, auditors, and compliance officers [25].

Data Governance and Security in Hybrid Credit Scoring

Hybrid cloud credit scoring systems have multi-layered security and governance controls. Some limitations still apply:

- **Data Residency Controls:** Sensitive credit data is kept in limited geographical locations, within on-premises or private cloud systems. Data that is not sensitive is either anonymized or aggregated and transferred to the public cloud for analysis [15].
- **Encryption Standards:** End-to-end encryption with AES-256 is applied for data both at rest and in transit. Encryption keys are separated for on-premises and cloud data and service components, which are managed by key management services [3].
- **Access Controls:** IAM frameworks provide role-based access control, but in a hybrid environment, consistent policy control cannot be assured [28].
- **Audit Trail Requirements:** In the Audit Trail Requirements, fragmented logging in hybrid environments leads to missing end-to-end auditability [28].

There is evidence that incorporating explainability into models is an important step to safeguard the integrity and improve compliance of models in AI/ML pipelines in hybrid cloud systems. However, contemporary models lack an explainability-governance framework [29].

Stakeholder Perspectives on Current Implementations

Survey results from banking executives indicate the following concerns related to the use of AI for credit scoring today [4]:

- 89% of banking institutions implement AI to perform real-time compliance checks, indicating a high level of use.
- Only 34% (according to the literature [4], [25], [11] have developed standardized explainability protocols aligned with regulatory standards.
- The issues faced include governance and explainability, regulatory ambiguity, and the computational burden of XAI methods [4], [20], [13]

We see a similar focus from Board members on the balance of operational benefit versus regulatory confidence. For AI/ML supported fraud deterrence, banking Board members see the operational benefit, though explainability and governance remain key concerns [4]. The results help refine the analysis of regulation to pinpoint the transparency requirements to be fulfilled via XAI.

US Regulatory Framework Analysis for AI Credit Scoring

This examines the US regulatory framework concerning the use of AI in credit scoring and the explainability and transparency requirements imposed by federal regulations.

Federal Reserve SR 11-7: Model Risk Management

The Federal Reserve's SR 11-7 provides extensive requirements on Model Risk Management (MRM) for credit scoring systems and related to the governance of Model's purposes, design and methods, limitations and the requirement to provide audit-ready explainability for the AI's input to decision [8]. The suggestions need to have their ideas validated independently, require performance to be continually assessed, be checked against other options, and be part of a centralized model inventory with clearly defined ownership and escalation processes [2], [5]

The current application of AI does not satisfy the standards of SR 11-7 since large language models and deep neural networks have little to no documentation that outlines the rationale behind their decisions beyond descriptions of their architecture [5]. Practical compliance is a challenge for banks: independent validation is difficult without specialized expertise, which is not always available [30], aggregation metrics in continuous monitoring do not explain individual decision changes [25], and hybrid cloud deployments make it hard to keep track of inventory in a distributed environment, and thus create lost audit trails [5].

Fair Housing Act (FHA) and Equal Credit Opportunity Act (ECOA)
FHA and ECOA ban credit decisions based on protected characteristics (race, color, religion, national origin, sex, marital status, age) and introduce the disparate impact doctrine, which bans facially neutral criteria that have a disparate impact [9]. ECOA requires explanations for credit denials to be "specific" (15 U.S.C. § 1691(d)), which must be understandable by humans and meet legal sufficiency requirements, and requires that AI systems provide explanations that would be understandable by humans and satisfy legal sufficiency requirements, and that analysis is performed to detect proxy discrimination when seemingly neutral features are correlated with protected features [11], [27]

The current post-hoc XAI methods (SHAP, LIME) generate technical feature scores that are not easy to understand by consumers and do not meet the requirement for providing a legally adequate adverse action explanation [13]. The lack of transparency in AI systems' credit decisions is also a key concern, as confirmed by research [5], and the proxy discrimination detection is still not complete, while the need for more human-friendly explanations remains a significant gap [13].

NIST AI Risk Management Framework (AI RMF)
According to NIST's AI RMF (January 2023) in order to implement AI systems, the financial services industry needs to have a high degree of trust and confidence in the systems which means that AI systems need to be transparent to the industry and stakeholders, and the outputs of the systems need to be interpretable and explainable (i.e. understand the outputs and know how to act on them) [10]. In the framework, credit scoring systems are in the "high risk" category and require additional controls for documenting the intended use of the credit scoring models, the instructions to build the models, the evaluation criteria and procedures, the evaluation results and limitations of the models, and the procedures to identify, quantify and mitigate bias in the models [10].

Existing hybrid cloud AI credit scoring systems suffer from systematic NIST AI RMF gaps: Model cards are incomplete [10], Bias management is ad hoc and not systematic [9], and Stakeholder communication focuses on technical audiences, not consumers and regulators [11]. Explanations to stakeholders are not adequate, with existing versions of the framework offering only technical metrics which are not appropriate for non-technical stakeholders like credit applicants, regulators, compliance officers, etc. [25].

Office of the Comptroller of the Currency (OCC) Fair Lending Standards
The OCC fair lending guidance expands anti-discrimination obligations to include operational expectations for AI credit systems, such as supporting examination procedures like comparative file review, statistical analysis of disparate treatment and disparate impact, and policy evaluation [9]. In addition to third-party risk management for cloud-deployed systems and ongoing monitoring of fairness metrics across demographic segments, AI credit scoring systems should also deliver explainable decision records, along with documentation of the model's contribution to the decision [30].

The individual decision explanations are still not adequate for the examination procedures of OCC and the monitoring of fairness metrics is inconsistent across banking institutions [9]. Research shows that in US banking, the OCC regulates the use of AI, and has shown concern with issues like bias and lack of transparency in credit approval models; currently, there are not sufficient tools or methods for detecting proxy discrimination, where facially neutral features correlate with protected characteristics, and there are not sufficient tools or methods for identifying proxy variables in complex feature interactions [30], [5].

Consumer Financial Protection Bureau (CFPB) Expectations
For AI credit decisions, regardless of model complexity, CFPB clarifies that the ECOA adverse action requirements continue to apply, including no "black box" exception (no exceptions for model

complexity), accuracy requirement (reasons must accurately reflect actual model decision factors), and the comprehensibility standard (reasons must be understandable to applicants without technical expertise) [11], [27]. The explanations generated by post-hoc explainability techniques should be accurate (faithful to the model's logic), specific (identifying key contributing factors) and comprehensible (accessible to non-technical consumers).

Existing "black box" AI models do not meet all three CFPB criteria, and the scores of post-hoc feature importance do not meet the legal sufficiency requirements for plain language explanations [13]. There is a considerable risk of enforcement due to the lack of specificity in explanations, as they are not actionable (e.g., "credit utilization: 0.23" instead of "Your credit card balances are 15% higher than approved applicants") and are not compared to other consumer groups (e.g., other consumers with similar credit card balances), and because they do not include information on remedial steps consumers must take to meet consumer protection standards [25], [11].

Regulatory Framework Evolution and Emerging Requirements

Transparency demands are evolving, including algorithmic accountability, the assessment of AI impacts and consumer rights to challenge AI-made decisions, with new bills proposed at different governance levels. The EU AI Act, as the first of its kind, has a risk-centric regulatory framework and creates new obligations to govern AI's application in banking systems. It compels the existence of technical documentation, audit trails, human oversight, and the monitoring of systems after deployment. These build requirements for convergence, while the regulatory framework for AI in financial services continues developing with predictions for explainability, and accountability and risk frameworks.

New proposed laws at the federal and state levels impose further transparency requirements, such as algorithmic accountability requirements, AI impact assessments for high-risk applications, and consumer rights to challenge AI decisions. The EU AI Act marks the first comprehensive AI legislation, introducing risk-based regulatory frameworks that directly affect the use of AI in banking institutions [31], mandates detailed technical documentation, audit trail features, human oversight, and post-market monitoring systems, driving convergence as the regulatory framework for AI in financial services is rapidly changing, with expectations of explainability, accountability and risk management [32].

The pressure on US banking institutions to implement strict transparency requirements to enable their operations in other countries is intensifying, but existing transparency practices do not have standardized explainability protocols in line with changing regulatory expectations. Hybrid cloud deployments introduce additional complexity through data residency requirements triggering cross-jurisdictional obligations [3], [15], third-party risk management challenges for cloud providers [5], fragmented audit trail continuity across distributed environments [28], and model governance complications in inventory, version control, and change management [5]. The future of AI in financial services will bring more challenges, including enhancing the transparency and auditability of AI models, tackling ethical issues, and developing standardized governance frameworks for AI in the financial industry [15].

Synthesis: Regulatory Transparency Requirements

In US banking, AI credit scoring systems must meet the following transparency requirements, which adds to the complexity of deploying on a hybrid cloud:

Table 4: Synthesized Regulatory Transparency Requirements

Regulatory Framework	Specific Transparency Requirements	Current Implementation Gap
SR 11-7 MRM	Comprehensive model documentation; validation protocols; ongoing	Deep learning models lack decision logic documentation; hybrid cloud deployments

	monitoring with decision-level audit trails	fragment audit trails
FHA/ECOA	Specific, accurate, comprehensible adverse action reasons; proxy discrimination detectability	Post-hoc XAI produces technical feature scores, not legally sufficient explanations; proxy discrimination detection incomplete
NIST AI RMF	Stakeholder-appropriate explanations; comprehensive model cards; systematic bias management	Model cards rarely complete; bias management ad hoc; explanations prioritize technical over consumer stakeholders
OCC Fair Lending	Support for comparative file review and statistical analysis; demographic outcome disaggregation	Individual decision explanations insufficient for examination procedures; fairness metrics monitoring inconsistent
CFPB ECOA Interpretation	Accurate, specific, comprehensible explanations regardless of model complexity	“Black box” models fail to produce explanations satisfying all three criteria simultaneously

IV. REGULATORY COMPLIANCE CHALLENGES IN HYBRID CLOUD ARCHITECTURES

Hybrid cloud deployment introduces additional regulatory complexity:

- Data Residency and Sovereignty: Cloud components can process data in different jurisdictions, thus introducing further compliance requirements [3], [15].
- Third-Party Risk Management: Cloud service providers are third-party relationships that must be managed and monitored, and contingency plans prepared [5].
- Audit Trail Continuity: Integrated logging infrastructure that is not commonly deployed in hybrid environments is required for Decision Provenance Tracking [28].
- Model Governance Fragmentation: Managing on-premises and cloud models using model inventory, version control, and change management makes it difficult to comply with SR 11-7 [5].

Researchers have highlighted the need for advancing transparency and auditability of AI models, tackling ethical issues, and developing standardized governance frameworks for AI in the financial industry for future work [15].

V. REGULATORY IMPERATIVE FOR EXPLAINABLE AI

Strong transparency, explainability, and auditability are essential in the regulatory framework for AI-based credit scoring in US banking. However, current implementations face several challenges to fulfill these requirements, such as lack of documentation of complex AI model decision logic, limited explanations produced by post-hoc XAI techniques, lack of thoroughness in verifying fairness in the presence of proxy discrimination, lack of audit trail in hybrid cloud environments, and explanations that are difficult for consumers to understand. The identified gaps indicate the need for fundamentally new explainable AI (XAI) frameworks for hybrid cloud banking systems. Here, explainability must be integrated as a core element rather than added as an optional feature after the fact. In the absence of these frameworks, the XAI for hybrid cloud banking must be built with explainability as a core element and not an optional feature.

Taxonomy of Explainability Gaps in Hybrid Cloud AI Credit Scoring

This section describes and categorizes the explainability and transparency gaps that prevent AI-based credit scoring systems in hybrid cloud banking from satisfying compliance with the formal requirements of US-based banking regulations regarding the use of AI.

Conceptual Framework for Gap Taxonomy
We classify the gaps in 3 dimensions:

- **Technical Dimension:** Gaps resulting from the limitations of AI model design and architectures, hybrid cloud infrastructure, and computational limitations.
- **Regulatory Dimension:** Gap between available solutions and specific regulatory requirements (SR 11-7, FHA/ECOA, NIST AI RMF).
- **Stakeholder Dimension:** Gaps resulting from lack of explainations to different stakeholder groups (end users, regulators, auditors, senior management, etc.).

This three-dimensional taxonomy enables accurate mapping of explainability gaps to remediation strategies in future XAI frameworks.

Technical Explainability Gaps

The most significant Technical Dimension gap is Model Architecture Opacity of credit scoring models built with deep neural networks and LLMs. This gap is the result of large scale, high-dimensional parameter spaces (i.e, deep learning), and low semantic interpretability of model transformations and representations of learned knowledge [7], [5]. This Gap clearly violates NIST RMF AI transparency and SR 11-7 documentation requirements.

These are exacerbated by current limitations of Post-Hoc XAI: SHAP suffers from computational complexity, making it unsuitable for real-time deployment [13]; LIME has been shown to have limitations in its ability to accurately capture model behavior in local approximations [11]; and attention mechanisms are not guaranteed to reflect feature importance [7]. Existing methods of explainability have been shown to be effective in providing general information about the behavior of the model, but further research is needed to create more granular and user-friendly explainable AI (XAI) modules that can be used by non-technical stakeholders [25].

Hybrid Cloud Provenance Complexity is the case where decision making is distributed among components trained in one environment, features engineered in another, and inferences made in yet another, with varying logging infrastructure [29], [15]. This is a high severity gap that will make it difficult to implement end-to-end audit trails necessary for SR 11-7 and OCC examination procedures.

Feature Attribution Instability yields inconsistent explanations when similar applications are given different explanations, even when they have similar scores [33], [11], Lack of Counterfactual Explanations involves systems rarely giving actionable guidance to consumers and support to regulators to examine decision boundaries [25], [11].

Regulatory Compliance Gaps

AI credit scoring models have a Documentation Deficiency with SR 11-7, which includes the lack of detailed documentation of the decision logic behind models; lack of validation of explainability mechanisms; lack of analysis of limitations for hybrid cloud deployment; and lack of model inventory integration across environments [30].

The current explanations generated by XAI models show evidence of ECOA Adverse Action Explanation Insufficiency in the form of feature importance scores that are not actionable, technical terms that are not comprehensible to consumers, lack of a contextual comparison, and no improvement guidance [25]. Cost-sensitive ML-based credit scoring reveals that class-weight adjustments affect the behaviour of the model and decision thresholds, leading to better detection and greater opacity, which must be documented in the model's output or in the explanations of the decision-making process in compliance with ECOA [26]. These deficiencies directly contradict the SR 11-7 and ECOA requirements and result in an inability to obtain regulatory approval and/or enforcement action against institutions.

There are NIST AI RMF Stakeholder Communication Gaps in current explanations, such as consumers are provided with technical scores rather than plain language explanations, regulators are provided with aggregate statistics, and auditors are provided with architecture descriptions rather than explanations that are ready for validation [11].

Stakeholder-Specific Explainability Gaps

Existing implementations show Consumer-Facing Explanation Gaps, such as providing feature importance scores instead of consumer-friendly explanations, lacking a comparative analysis with accepted applicants, lack of actionable improvement suggestions, and technical presentations that diminish consumer trust [11], [13]. Lenders who focus on SME credit often complain about the challenge of maintaining a high level of predictive accuracy while providing clear and easily understood explanations to regulators of the reasons behind the risk decisions [34]. This gap is a direct failure of legal compliance for adverse action notice requirements under ECOA, and under CFPB comprehensibility standards.

Existing hybrid cloud AI deployments do not offer a complete audit trail for regulatory examiners to verify decisions. Current systems offer only high-level architectural descriptions, rather than decision logic specifications, limit access to training data because of hybrid cloud confidentiality constraints, and do not have standardized methods for evaluating the accuracy of explainability mechanisms [25], [30], and for making it easier for OCC examinations and SR 11-7 monitoring to verify. The need for interpretable and explainable AI models is highlighted, as it promotes transparency in decision making and builds trust between auditors and regulators [25].

While current technical explainability tools offer only technical metrics (model accuracy, SHAP values), and research shows that governance and explainability are still a concern for banking board members despite its potential operational benefits, it is clear that there are gaps in the board-level explainability necessary to meet SR 11-7 governance expectations and this is a medium-severity gap in governance quality, not a compliance issue [4].

Hybrid Cloud-Specific Explainability Gaps

Model decision provenance is split between on-premises, private cloud, and public cloud elements, making it difficult to create a unified audit trail necessary for SR 11-7 ongoing monitoring requirements and OCC examination procedures. Research highlights the need for embedding explainability to ensure integrity and compliance readiness of models, and future research should prioritize enhancing the transparency and auditability of AI models in hybrid settings [29], [15]. The use of cloud provider ML services (AWS SageMaker, Google Vertex AI, Azure ML) adds a layer of vendor opacity. According to BIS, while AI models can improve efficiency in AML/CFT compliance and the estimation of regulatory capital, the models' opacity and dependence on third parties lead to model and concentration risks, which can be mitigated through enhanced supervisory frameworks [5].

Prioritized Gap Taxonomy

Table 5: Prioritized Explainability Gap Taxonomy

Gap Category	Specific Gap	Severity	Primary Regulatory Impact	Required XAI Framework Capability
Technical	Model Architecture Opacity	Critical	SR 11-7, NIST AI RMF	Inherent interpretability or high-fidelity post-hoc explainability
Regulatory	ECOA Adverse Action Insufficiency	Critical	ECOA, CFPB	Consumer-comprehensible, legally sufficient explanations
Regulatory	FHA/ECOA Proxy Discrimination	Critical	FHA, ECOA, OCC	Systematic bias detection and feature governance
Regulatory	SR 11-7 Documentation Deficiency	Critical	SR 11-7	Comprehensive model documentation generation
Stakeholder	Consumer Explanation Gap	Critical	ECOA, CFPB	Plain-language, actionable explanations
Stakeholder	Regulator Audit Trail Gap	Critical	SR 11-7, OCC	End-to-end decision provenance tracking
Technical	Post-Hoc XAI Limitations	High	SR 11-7, CFPB	Validated, accurate explanation mechanisms

Gap Category	Specific Gap	Severity	Primary Regulatory Impact	Required XAI Framework Capability
Technical	Hybrid Cloud Provenance Complexity	High	SR 11-7, OCC	Unified audit trail across hybrid environments
Regulatory	Model Validation Gap	High	SR 11-7	Explainability validation protocols
Regulatory	NIST Stakeholder Communication	High	NIST AI RMF	Multi-stakeholder explanation generation
Hybrid Cloud	Cross-Environment Fragmentation	High	SR 11-7, OCC	Integrated logging and provenance tracking
Hybrid Cloud	Cloud Vendor Dependency	High	SR 11-7	Vendor-agnostic explainability frameworks
Stakeholder	Auditor Validation Gap	High	SR 11-7	Validator-friendly documentation and tools
Technical	Feature Attribution Instability	Medium	CFPB	Robust, stable explanation generation
Technical	Lack of Counterfactuals	Medium	ECOA, OCC	Counterfactual explanation capabilities
Stakeholder	Developer Debugging Gap	Medium	Internal governance	Fine-grained diagnostic explainability
Stakeholder	Executive Risk Communication	Medium	SR 11-7 governance	Business-oriented risk metrics
Hybrid Cloud	Data Sovereignty Constraints	Medium	Data residency regs	Localized explanation generation
Hybrid Cloud	Model Version Control	Medium	SR 11-7	Unified model registry and versioning

VI. DESIGN IMPLICATIONS FOR REGULATORY-ALIGNED EXPLAINABLE AI

In this section, the findings of the existing implementations review (Section III), the regulatory perspective (Section IV), and the gap analysis (Section V) are integrated to derive some design implications for constructing regulatory compliant explainable AI frameworks concerning hybrid cloud credit scoring within the US banking system. The evaluative findings suggest a design focus on the following four strategic priorities relating to the deployment of explainable AI systems:

Regulatory Compliance as Foundational Requirement

The absence of explainability should not be justified using the post hoc rationale for opt-in/opt-out regulatory compliance. Studies suggest that explainability should not be tacked on to the technical details of AI systems, but should be viewed as an integral component of responsible AI [27]. As part of the production process, fundamental transparency is a requirement for the comprehensive US regulatory framework (SR 11-7, FHA/ECOA, NIST AI RMF). XAI systems should go beyond the inference phase of the operational model lifecycle and incorporate the capability to interpret at all other phases of data prep, feature eng, model training and validation, deployment, and ongoing system oversight.

Hybrid Cloud Architecture as Design Constraint

Where deployment is in a hybrid cloud, the transparency challenges of vendor locked systems and fragmented provenance are beyond the scope of most XAI methods. Research indicates that transparent financial AI is attainable, but successful application requires a hybrid cloud-native architecture and development of bespoke explainability techniques [35]. The XAI frameworks need to be designed specifically for hybrid architectures, that can create a unified audit trail, track provenance across environments and include vendor-agnostic explainability components that ensure transparency even when deployed across multiple environments.

Multi-Stakeholder Explanation Requirements

There are different types of explanation that different stakeholders need. Research shows that detailed analysis of explainability needs of six stakeholder groups, most of which are nontechnical users, revealed that a single explanation format is not enough to satisfy the needs of stakeholders [11]. XAI frameworks need to create multiple artifacts of explanations that can be generated from a single set of model decision information, including consumer understandable adverse action notices, regulator accessible audit trails, auditor friendly validation documentation, developer focused debugging information, and executive level risk assessments from the same underlying model decision.

Explanation Quality Assurance

Existing post-hoc XAI methods (SHAP, LIME) do not provide a mechanism for measuring the accuracy of the explanations, that is, whether they are accurate in capturing the model's decision-making process. As the research points out, creating strong XAI evaluation guidelines will be a key factor in the responsible use of AI technologies [33]. Explanation validation mechanisms in XAI frameworks should cover three criteria: fidelity (does the explanation match the behavior of the model?), stability (do similar inputs to the model lead to similar explanations?), and completeness (are all material decision factors captured?).

VII. KEY FINDINGS

This research is the first in-depth study of the implementation of AI credit scoring in hybrid cloud banking systems, systematically identifying gaps in regulatory transparency that hinder compliance with banking standards in the United States.

Widespread Adoption with Persistent Opacity

89% of U.S. banking institutions are using AI to monitor compliance in real time, and it has become production-critical in these institutions [4]. The leap from simple machine learning to advanced deep learning and LLM systems ushers in a new decade of basic opacity that safeguards most transparency requirements. For contractual validation, the provision of interpretable models is always preferred over black-box models because black-box models will always make compliance verification extremely difficult [2].

Hybrid Cloud Deployment Fragmentation

Credit scoring solutions of the present have hybrid clouds that integrate both proprietary and public clouds to improve computational scalability and cost efficiency. However, the merging of cloud systems leads to the fragmentation of model decision provenance, presenting significant data residency, auditability and explainability challenges that the Bank for International Settlements has identified as key risk factors [5].

Regulatory-Technical Misalignment

U.S. banking regulations (SR 11-7 Model Risk Management, FHA/ECOA Fair Lending, NIST AI RMF) have transparent, explainable, auditable requirements. Existing AI solutions fall short of these requirements in many ways including:

- **Documentation Deficits:** Per SR 11-7, most Complex AI Systems do not document the decision logic sufficiently.
- **Explanation Inadequacy:** Post-hoc XAI techniques (SHAP, LIME) do not provide legally sufficient adverse action explanations in ECOA.
- **Fairness Verification Gaps:** Proxy discrimination detection and disparate impact assessment is still not complete.
- **Audit Trail Fragmentation:** Hybrid cloud architectures block end-to-end decision provenance for regulatory examinations – Audit Trail Fragmentation.

Multi-Stakeholder Explanation Deficits

While various stakeholders have very different needs for the type of explanation, the available implementations offer mainly technical feature importance scores, which are not suitable for most stakeholders. Existing explainability methods provide basic insights into model behavior, but future work must focus on developing more granular, user-friendly explainable AI (XAI) modules that cater to non-technical stakeholders such as regulators, auditors, and compliance officers [25], [12]

Critical vs. Addressable Gaps

Table 6: Classification of explainability gaps in AI credit scoring systems based on severity and regulatory compliance impact.

Gap Category	Explainability Deficits Identified
Serious Gaps (requiring immediate remediation)	Model architecture opacity; ECOA adverse action insufficiency; FHA/ECOA proxy discrimination detection limitations; SR 11-7 documentation deficiency; inadequate explanations for consumers and regulators
High-Priority Gaps (significantly impeding compliance)	Limitations of post-hoc XAI methods; hybrid cloud provenance complexity; model validation challenges; stakeholder communication gaps identified by NIST frameworks
Medium-Priority Gaps (enhancing but not preventing compliance)	Feature attribution instability; absence of counterfactual explanations; limitations in developer debugging capabilities; gaps in executive-level risk communication

VIII. RESEARCH CONTRIBUTIONS AND NOVELTY

This research makes notable contributions to explainable AI (XAI), compliance with regulation, and hybrid cloud banking. This research systematically investigates explainability gaps for hybrid cloud AI credit scoring, aligned with US regulations, for the first time, and empirically places them in order of importance. This study maps SR 11-7, FHA/ECOA, and NIST's AI RMF to technical explainability. This research provides a means to understand the explanation needs of different stakeholders (i.e., consumers, regulators, auditors, developers, executives), as opposed to a one-size-fits-all approach. This research proposes the design principles of XAI for hybrid cloud with provenance, vendor lock independence, and distributed audit trails. This research empirically maps current AI credit scoring models in financial services and provision constraining deployment.

IX. FUTURE RESEARCH DIRECTIONS

Research in the area of regulation-compliant systems for hybrid cloud banking with explainable AI may take the following routes, using this study as a starting point:

1. Credit scoring based on ML in hybrid cloud banking is influenced significantly by the financial attributes of different classes of borrowers (i.e., consumers, micro and small enterprises) regardless of formulation. Future research should thus focus on segment-adaptive XAI frameworks that would adapt explanations for each segment's specific credit-risk profile [36].
2. Developing benchmarks that are accepted by the regulatory community for explainability quality, fidelity, and completion the creation of robust XAI evaluation standards will be a key factor in the responsible use of AI technologies [33].
3. Banking institutions' investigation of hybrid explainability models (human-AI interaction) may provide valuable insights into their effectiveness and user acceptance [27].
4. Research on improving computational efficiency in real-time financial applications, for low-latency generation of explanations without compromising quality [11].
5. Risk and defense mechanisms of explanation manipulation and explanation integrity.
6. Exploration of generalizability of credit scoring XAI frameworks to other banking AI applications such as fraud detection, AML monitoring, and investment advisory.

X. CONCLUSION

The use of AI in credit scoring has many advantages, such as better risk assessment, financial inclusion with alternative data, and operational efficiency. However, the benefits of these come with the need for high explainability to ensure regulatory adherence and consumer confidence. The current AI credit scoring systems in hybrid cloud banking environments reveal significant transparency deficits, which are incompatible with the regulatory standards enforced by the USA, and may thus pose a compliance risk and damage the reputation of banking institutions. Purpose-built XAI frameworks are required to address these gaps. This study provides a foundational framework to build XAI systems on hybrid cloud banking that are compliant with regulations and focused on providing solutions.

REFERENCES

1. A. Sirangi, "AI-Driven Risk Scoring Engine for Financial Compliance in Multi-Cloud Environments," Journal of Electrical Systems, vol. 17, no. 1, pp. 138–150, Jan. 2021, doi: 10.52783/JES.8887.
2. T. M. Chettier, V. A. Kumar Boyina, and S. Rangineni, "AI-Powered Risk Assessment and Compliance in Cloud Cybersecurity," International Journal of Computer Trends and Technology, vol. 73, no. 3, pp. 49–56, Mar. 2025, doi: 10.14445/22312803/IJCTT-V73I3P107.
3. "Grappling With Financial Services Regulation? Here's How AI Can Help You Comply."
4. A. Eskandarany, "Adoption of artificial intelligence and machine learning in banking systems: a qualitative survey of board of directors," Front. Artif. Intell., vol. 7, 2024, doi: 10.3389/FRAI.2024.1440051.
5. J. C. Crisanto, C. B. Leuterio, J. Prenio, and J. Yong, "Regulating AI in the financial sector: recent developments and main challenges," Dec. 2024, [Online]. Available: https://www.bis.org/fsi/publ/insights63.htm
6. Y. Wang and X. Yang, "Machine Learning-Based Cloud Computing Compliance Process Automation."
7. M. Golec and M. AlabdulJalil, "Interpretable LLMs for Credit Risk: A Systematic Review and Taxonomy," Jun. 2025, Accessed: Jan. 11, 2026. [Online]. Available: https://arxiv.org/pdf/2506.04290
8. F. Reserve Board, "What Are We Learning about Artificial Intelligence in Financial Services?," 2018. [Online]. Available: http://www.fsb.org/wp-content/uploads/P011117.pdf.
9. G. Valdrighi et al., "Best Practices for Responsible Machine Learning in Credit Scoring," Neural Comput. Appl., vol. 37, no. 25, pp. 20781–20821, Sep. 2024, doi: 10.1007/s00521-025-11520-y.
10. N. Institute of Standards, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," Jan. 2023, doi: 10.6028/NIST.AI.100-1.
11. C.-A. Wilson, "Explainable AI in Finance: Addressing the Needs of Diverse Stakeholders," Aug. 2025, doi: 10.56227/25.1.25.
12. M. Bücker, G. Szepannek, A. Gosiewska, and P. Biecek, "Transparency, Auditability and eXplainability of Machine Learning Models in Credit Scoring," Journal of the Operational Research Society, vol. 73, no. 1, pp. 70–90, Sep. 2020, doi: 10.1080/01605682.2021.1922098.
13. A. John and A. Ahsun, "Explainable AI (XAI) for Fraud Detection: Building Trust and Transparency in AI-Driven Financial Security Systems Authors," Jun. 2025, doi: 10.2139/SSRN.5285281.
14. M. K. Nallakaruppan, B. Balusamy, M. L. Shri, V. Malathi, and S. Bhattacharyya, "An Explainable Ai Framework for Credit Evaluation And Analysis," 2023, doi: 10.2139/SSRN.4543661.
15. K. Sankaralingam, "Cloud AI for Regulatory Compliance and Risk Management in Financial Institutions," International Journal of Computing and Engineering, vol. 7, no. 4, pp. 8–15, May 2025, doi: 10.47941/IJCE.2759.
16. Narendra Bhargav Boggarapu, "Modernizing Banking Compliance: An Analysis of AI-Powered Data Governance in a Hybrid Cloud Environment," International Journal of Scientific Research in Computer Science, Engineering and Information Technology, vol. 10, no. 6, pp. 2373–2381, Dec. 2024, doi: 10.32628/CSEIT2410612434.
17. L. G. Anand, "AI-Powered Regulatory Technologies Transforming Compliance in U.S. Financial Institutions," Jun. 2025, doi: 10.2139/SSRN.5349014.
18. "Artificial Intelligence in Financial Services," 2025.
19. "Generative AI in Banking-A Comprehensive Overview What do banks need to consider before jumping into the rabbit hole? Expert report of the SBA," 2025.

20. K. Goyal, M. Garg, and S. Malik, "Adoption of artificial intelligence-based credit risk assessment and fraud detection in the banking services: a hybrid approach (SEM-ANN)," Future Business Journal 2025 11:1, vol. 11, no. 1, pp. 1–20, Mar. 2025, doi: 10.1186/S43093-025-00464-3.
21. I. Iaroshev et al., "Evaluating Retrieval-Augmented Generation Models for Financial Report Question and Answering," Applied Sciences 2024, Vol. 14, vol. 14, no. 20, Oct. 2024, doi: 10.3390/APP14209318.
22. Y. Gao et al., "Retrieval-Augmented Generation for Large Language Models: A Survey," Proceedings - 2024 Conference on AI, Science, Engineering, and Technology, AIxSET 2024, pp. 166–169, Dec. 2023, doi: 10.1109/AIxSET62544.2024.00030.
23. S. Farquhar, J. Kossen, L. Kuhn, and Y. Gal, "Detecting hallucinations in large language models using semantic entropy," Nature 2024 630:8017, vol. 630, no. 8017, pp. 625–630, Jun. 2024, doi: 10.1038/s41586-024-07421-0.
24. A. Ali et al., "Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review," Oct. 01, 2022, MDPI. doi: 10.3390/app12199637.
25. V. K. Tambi, "Real-Time Compliance Monitoring in Banking Operations using AI," Jan. 01, 2022. [Online]. Available: https://papers.ssrn.com/abstract=5366168
26. B. Hover, "Applying Cost-Sensitive Machine Learning Models to Loan Default Prediction," 2022.
27. A. Lindholm and A. Cederörn, "Unveiling the Black Box: Explainability and trust in AI-based banking systems A systematic literature review on transparency challenges in autonomous financial decision-making."
28. K. Torkura, M. I. H. Sukmana, F. Cheng, and C. Meinel, "Continuous Auditing & Threat Detection in Multi-Cloud Infrastructure," Oct. 2020, doi: 10.36227/TECHRXIV.13108313.
29. S. R. Sudharsanam, D. Venkatachalam, and D. Paul, "Securing AI/ML Operations in Multi-Cloud Environments: Best Practices for Data Privacy, Model Integrity, and Regulatory Compliance."
30. J. Kokina, S. Blanchette, T. H. Davenport, and D. Pachamanova, "Challenges and opportunities for artificial intelligence in auditing: Evidence from the field," International Journal of Accounting Information Systems, vol. 56, p. 100734, Dec. 2025, doi: 10.1016/J.ACCINF.2025.100734.
31. M. Mucsková, "Transforming Banking with Artificial Intelligence," Trends Economics and Management, vol. 18, no. 42, pp. 21–37, Aug. 2024, doi: 10.13164/TRENDS.2024.42.21.
32. "Global Financial Services: The Sector's Current AI Regulatory Landscape – Publications." [Online]. Available: https://www.morganlewis.com/pubs/2024/05/global-financial-services-the-sectors-current-ai-regulatory-landscape
33. P. Seth and V. K. Sankarapu, "Bridging the Gap in XAI-Why Reliable Metrics Matter for Explainability and Compliance," Feb. 2025, [Online]. Available: https://arxiv.org/pdf/2502.04695
34. E. Bickmore, A. MacKinlay, and Y. Tellez, "Old Program, New Banks: Online Banks in Small Business Lending," SSRN Electronic Journal, Mar. 2024, doi: 10.2139/SSRN.4537337.
35. S. K. Cherukuru, "Explainable AI (XAI) in Cloud-Native Financial Services: Building Trust and Transparency in Modernized Decision Engines," 2025, doi: 10.32996/jcsts.
36. T. T. Yussuph, "Leveraging Machine Learning Algorithm to Enable Access to Credit for Small Businesses in the United States of America," International Journal on Cybernetics & Informatics, vol. 13, no. 1, pp. 53–65, Feb. 2024, doi: 10.5121/ijci.2024.130105.