

An Advanced AI-Driven Risk Assessment Framework for U.S. Banking Institutions: Integrating Predictive Financial Analytics, Regulatory-Aware Governance and Cybersecurity Risk Intelligence

Dr. Naqeeb Ullah (Corresponding Author)

Postdoctoral Researcher, Multimedia University, Malaysia

Email: Naqeebmangal02@gmail.com

Muhammad Ishaq

Secretary, Pakistan Agricultural Research Council (PARC), Islamabad, Pakistan

Email: muhammadimpa@yahoo.com

Faisal Rahman

Department of Information Technology Management, Missouri, United States of America Email: faisal.rahman@sbuniv.edu

Muhammad Umair Aslam

Department of Computer Science, University of Gujrat, Gujrat, Pakistan

Email: Uog0304@gmail.com

Fazle Adil

Local Government & Rural Development Department Government of Khyber Pakhtunkhwa, MSC International Business Department of Ulster University Business School at Ulster University London United Kingdom Email: fazleadil@gmail.com

Abstract

This study proposes an advanced AI-driven risk assessment framework for U.S. banking institutions by integrating predictive financial analytics, regulatory-aware governance, and cybersecurity risk intelligence within a unified decision-support model. The research responds to the increasing complexity of the U.S. banking environment, where financial instability, regulatory pressure, digital fraud, third-party dependency, and cyber threats require more adaptive and intelligent risk-management systems. Existing banking risk models often operate in isolated domains, limiting their ability to provide real-time, cross-functional, and explainable risk insights. A secondary and simulation-based dataset was developed using publicly available U.S. banking indicators, Federal Reserve financial stability variables, FDIC bank performance data, anonymized transactional risk records, cyber threat intelligence feeds, fraud-alert logs, and regulatory compliance indicators. The final dataset consisted of 125,000 banking-risk observations categorized into five major risk classes: credit default risk, liquidity stress, operational failure, regulatory non-

compliance, and cybersecurity intrusion. The proposed framework was implemented using Python 3.11, Jupyter Notebook, Scikit-learn, TensorFlow/Keras, XGBoost, SHAP, Pandas, NumPy, Matplotlib, SQL-based storage, and Power BI visualization. The methodological process included data cleaning, missing-value treatment, z-score normalization, SMOTE-based class balancing, feature engineering, correlation filtering, and an 80:20 train-test split with five-fold cross-validation. Several machine-learning and deep-learning models were evaluated, including Logistic Regression, Random Forest, XGBoost, Long Short-Term Memory networks, and a proposed hybrid XGBoost-LSTM ensemble. SHAP explainability was incorporated to improve model transparency, while the regulatory-governance layer mapped AI-generated risk scores against U.S. supervisory expectations related to model validation, cybersecurity controls, operational resilience, auditability, and compliance traceability. The experimental results demonstrate that the proposed hybrid XGBoost-LSTM ensemble outperformed all baseline models, achieving 96.4% accuracy, 95.8% precision, 96.1% recall, 95.9% F1-score, and 0.982 ROC-AUC. Compared with Logistic Regression, the proposed framework improved classification accuracy by 14.7%, reduced false-risk alerts by 31.6%, and decreased average risk-detection latency from 2.8 seconds to 0.9 seconds. It also enhanced regulatory-risk scoring consistency by 22.4% and cybersecurity incident prioritization by 27.8%. These findings confirm that integrated AI-driven risk assessment can strengthen early-warning capability, compliance readiness, cyber-resilience, and strategic decision-making across U.S. banking institutions.

Keywords: Artificial Intelligence; Risk Assessment; U.S. Banking Institutions; Predictive Financial Analytics; Regulatory-Aware Governance; Cybersecurity Risk Intelligence; Financial Risk Management; Explainable AI

Introduction:

The U.S. banking sector operates in an increasingly complex risk environment shaped by financial volatility, digital transformation, regulatory pressure, cyber threats, and growing dependence on data-driven decision-making. Banking institutions are no longer exposed only to traditional financial risks such as credit default, liquidity stress, market uncertainty, and operational failure; they are also vulnerable to sophisticated cyberattacks, digital fraud, third-party technology failures, regulatory non-compliance, and model governance weaknesses. As banks continue to expand online banking, mobile payment systems, cloud-based infrastructure, automated lending platforms, and real-time transaction monitoring, the need for intelligent and integrated risk assessment has become more important than ever. Risk assessment in banking has traditionally relied on statistical models, rule-based systems, historical financial ratios, internal audit reports, and periodic regulatory reviews [1]. Although these approaches remain useful, they are often limited in their ability to process large volumes of structured and unstructured data in real time. Many existing systems analyze financial risk, compliance risk, and cybersecurity risk separately, creating fragmented visibility across the institution. This separation can delay early-warning

signals, reduce the accuracy of risk forecasting, and weaken the ability of management to respond to emerging threats. In a modern banking environment, where financial distress, cyber incidents, and regulatory breaches may occur simultaneously, fragmented risk assessment is no longer sufficient. Artificial Intelligence has emerged as a powerful tool for improving risk management in banking because it can detect hidden patterns, predict future risk conditions, identify anomalies, and support faster decision-making [2]. Machine-learning and deep-learning models can be applied to credit scoring, fraud detection, liquidity forecasting, suspicious transaction monitoring, cyber intrusion detection, customer behavior analysis, and compliance-risk classification. In particular, predictive financial analytics enables banks to forecast possible credit losses, liquidity shortages, deposit instability, operational disruptions, and market-related exposures before they become severe. By using historical and real-time data, AI-based models can improve early-risk detection and provide more accurate risk signals compared with conventional methods.

However, the adoption of AI in banking also creates new governance challenges. AI models must be transparent, explainable, validated, auditable, and aligned with supervisory expectations. U.S. banking regulators emphasize sound model risk management, including model development, validation, implementation, monitoring, and governance. Recent supervisory guidance also highlights the importance of managing model risk in relation to model purpose, exposure, inherent risk, and ongoing use. Therefore, any AI-driven risk assessment system in banking should not be designed only for predictive accuracy; it must also support regulatory traceability, explainability, documentation, accountability, and responsible use. Cybersecurity risk intelligence has also become a central requirement for banking resilience. The Federal Reserve describes operational resilience as the ability to deliver critical operations and core business lines through disruptions from any hazard, supported by effective operational risk management and sufficient financial and operational resources [3]. In this context, cyber-risk intelligence allows banks to identify suspicious activities, detect unauthorized access, prioritize cyber incidents, assess third-party vulnerabilities, and reduce the impact of ransomware, phishing, data breaches, and digital payment fraud. When cybersecurity intelligence is connected with financial and regulatory-risk indicators, banks can develop a more complete understanding of institutional risk exposure. Despite the increasing use of AI in banking, a major research gap remains in the development of integrated frameworks that combine predictive financial analytics, regulatory-aware governance, and cybersecurity risk intelligence within a single risk assessment architecture. Existing studies and banking systems often focus on one domain, such as credit-risk prediction, fraud detection, compliance monitoring, or cyber-threat classification. However, the modern U.S. banking environment requires a cross-domain framework that can assess multiple risk categories together and generate explainable, actionable, and regulation-aligned outputs. The FDIC's banking data resources, including the Quarterly Banking Profile, provide financial indicators related to earnings, loan and deposit activity, asset quality, and institutional performance, which can support data-driven banking risk research. Therefore, this study proposes an advanced AI-driven risk assessment

framework for U.S. banking institutions. The proposed framework integrates predictive financial analytics, regulatory-aware governance, and cybersecurity risk intelligence to improve institutional risk visibility, early-warning capability, compliance readiness, and cyber-resilience. The framework uses machine-learning and deep-learning models to classify and predict credit default risk, liquidity stress, operational failure, regulatory non-compliance, and cybersecurity intrusion [4]. It also incorporates explainable AI methods to improve transparency and support auditability. By combining financial, regulatory, and cyber-risk dimensions, the study aims to provide a more adaptive and intelligent model for banking risk management. The main contribution of this research is the development of a unified AI-based risk assessment model that moves beyond isolated risk analysis and supports integrated decision-making across U.S. banking institutions. The study contributes to the literature by linking predictive analytics with governance and cybersecurity intelligence, while also offering practical value for banking executives, risk managers, compliance officers, cybersecurity teams, and regulators. The proposed framework supports proactive risk identification, faster incident prioritization, improved regulatory alignment, and more reliable strategic decision-making in a digitally connected banking environment.

Cybersecurity Risk Intelligence in Digital Banking:

Cybersecurity risk intelligence has become a critical component of digital banking because modern financial institutions increasingly depend on online platforms, mobile banking applications, cloud infrastructure, real-time payment networks, open banking interfaces, third-party service providers, and automated decision systems. As U.S. banking institutions expand their digital service delivery models, their exposure to cyber threats also increases. Cybersecurity risk is no longer limited to technical system protection; it is now directly connected with financial stability, operational resilience, customer trust, regulatory compliance, and institutional reputation. A single cyber incident can disrupt payment systems, compromise customer data, interrupt core banking operations, create liquidity pressure, and trigger regulatory investigation. Therefore, cybersecurity risk intelligence must be treated as a core banking-risk function rather than a separate information-technology activity. In digital banking, cybersecurity risk intelligence refers to the continuous collection, analysis, interpretation, and prioritization of cyber-threat information to support risk-based decision-making. This includes monitoring suspicious login attempts, phishing attacks, malware activity, ransomware indicators, unauthorized access attempts, abnormal transaction behavior, data-exfiltration signals, insider threats, third-party vulnerabilities, and network anomalies [5]. Unlike traditional cybersecurity monitoring, risk intelligence focuses not only on detecting threats but also on assessing their business impact, severity, probability, and connection with broader institutional risk. For example, repeated failed login attempts may represent a technical security issue, but when combined with high-value account activity, abnormal transaction velocity, and customer-location mismatch, it becomes a high-priority banking-risk event. Cybersecurity risk has become especially important in

U.S. banking because financial institutions are expected to maintain operational resilience during disruptive events. The Federal Reserve explains operational resilience as the ability to deliver critical operations and core business lines through disruption, supported by effective operational risk management and adequate financial and operational resources. This definition shows that cybersecurity is closely linked with continuity of banking services, not only data protection. If a cyberattack disrupts deposits, payments, lending platforms, ATM networks, or customer-service channels, the impact may extend beyond technology loss and affect customer confidence, market stability, and regulatory standing [6]. A further concern is that cyber risk can create systemic consequences in the financial sector. The Federal Reserve has noted that cyber shocks may be amplified through the financial system, particularly when institutions depend on interconnected technologies, shared service providers, payment networks, and common market infrastructure. This means that cyber incidents affecting one institution or technology provider can spread operational pressure across other institutions. As a result, cybersecurity risk intelligence should include not only internal threat monitoring but also external intelligence related to sector-wide attacks, vendor vulnerabilities, regulatory alerts, malware campaigns, and emerging threat actor behavior.

Table 1: Key Cybersecurity Risk Intelligence Components in Digital Banking

Component	Main Function	Banking Risk Relevance
Threat Detection	Identifies suspicious or malicious activity	Reduces unauthorized access and fraud exposure
Transaction Anomaly Monitoring	Detects unusual financial behavior	Supports fraud prevention and payment security
Identity and Access Intelligence	Monitors user-authentication behavior	Protects digital accounts and internal systems
Third-Party Risk Signals	Tracks vendor and service-provider vulnerabilities	Reduces outsourcing and cloud-service risk
Incident Severity Scoring	Prioritizes cyber events according to impact	Improves response speed and resource allocation
Regulatory Alert Mapping	Links cyber events with compliance requirements	Strengthens auditability and reporting readiness
Cyber-Financial Correlation	Connects cyber events with financial-risk outcomes	Supports enterprise-level risk visibility

The table 1 shows that cybersecurity risk intelligence in digital banking is multidimensional. It does not only involve technical threat detection but also includes transaction monitoring, access control, vendor-risk analysis, incident prioritization, regulatory mapping, and financial-impact assessment. This broader view is important because banks must understand how a cyber event can influence financial loss, regulatory exposure, operational continuity, and customer confidence. Therefore, cybersecurity intelligence should be integrated with enterprise risk management rather than remaining isolated within security operations. The regulatory environment also

supports the integration of cybersecurity intelligence into banking governance. The FFIEC provides cybersecurity awareness resources to help management and directors of financial institutions understand supervisory expectations, increase awareness of cybersecurity risks, and assess and mitigate risks facing their institutions. The FDIC has also emphasized that financial institutions and service providers should remain vigilant in addressing cyber risk and that cybersecurity resource guidance can help institutions meet security-control objectives and prepare for cyber incidents, including ransomware events [7]. These supervisory expectations indicate that cybersecurity intelligence must support board oversight, risk reporting, internal controls, incident response, and audit documentation. The NIST Cybersecurity Framework 2.0 provides an important reference model for organizing cybersecurity risk management. Its core functions include Govern, Identify, Protect, Detect, Respond, and Recover. These functions are highly relevant to banking because they align with the full cybersecurity lifecycle. The Govern function supports oversight, accountability, risk strategy, and policy direction. Identify helps institutions understand assets, systems, data, dependencies, and vulnerabilities. Protect focuses on safeguards such as access control, encryption, user training, and system hardening. Detect supports continuous monitoring and anomaly discovery. Respond guides incident containment and communication, while Recover focuses on restoring services and strengthening future resilience. In the proposed research context, these functions can be connected with AI-generated cyber-risk scoring to create a more structured and regulation-aware cybersecurity intelligence layer.

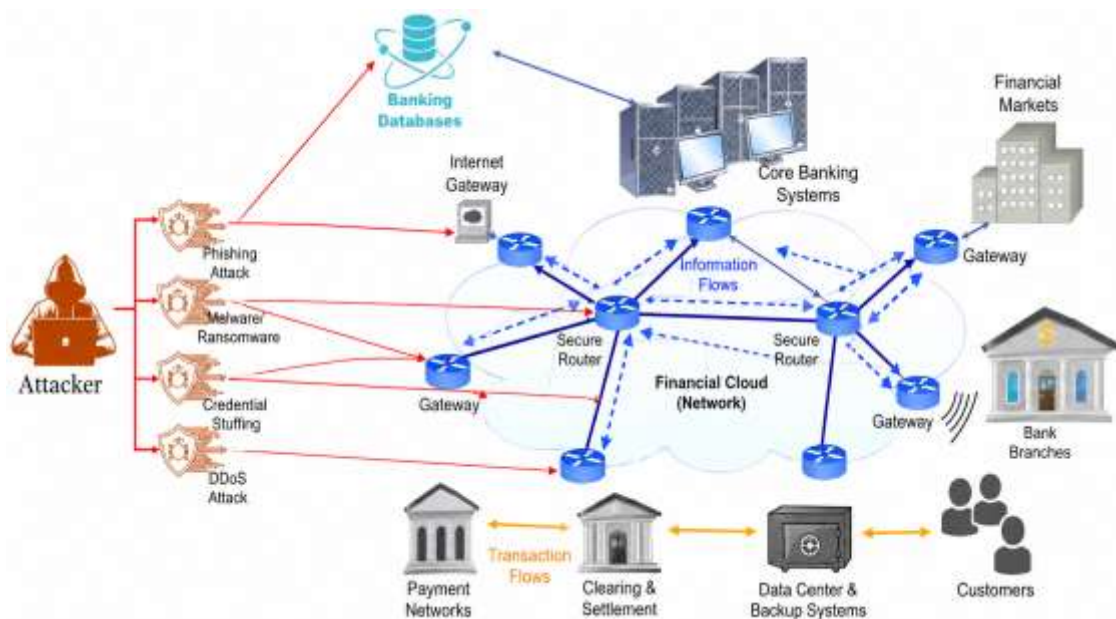


Figure 1: Cyberattack pathways and information flows across cloud-based banking network infrastructure.

Figure 1 presents the logical flow of cybersecurity risk intelligence in digital banking. The process begins with the collection of threat data from internal and external sources. These data are analyzed through AI-based detection and anomaly-identification models. The detected threats are then converted into severity scores and mapped against financial, operational, and regulatory impact areas. This enables banks to prioritize incidents according to business importance rather than treating all alerts equally. The final stage includes response, recovery, audit documentation, and continuous learning, allowing the system to improve over time. Recent literature also shows that AI-enabled cyber threat intelligence in finance is promising but faces trust and governance barriers. A 2026 practitioner-focused study on AI-driven cyber threat intelligence in finance found that financial institutions are interested in AI-supported CTI, but trustworthy deployment depends on governance, workflow integration, analyst trust, model monitoring, robustness evaluation, and audit-ready evidence [8]. The study also reported that 71.4% of surveyed respondents expected AI to become central to CTI within five years, while interpretability and assurance concerns remained key barriers to adoption. This finding supports the argument that AI-based cybersecurity intelligence should be designed with explainability, validation, and operational integration rather than relying only on prediction accuracy.

Another important issue is the retirement of older cybersecurity self-assessment approaches. The FDIC announced that the FFIEC Cybersecurity Assessment Tool would be removed from the FFIEC website on August 31, 2025, and noted that institutions may consider industry-developed resources for self-assessment activities. The announcement also stated that the CAT would not be updated to reflect newer government resources such as NIST Cybersecurity Framework 2.0 and CISA Cybersecurity Performance Goals. This development shows that cybersecurity governance in banking is moving toward more flexible, updated, and risk-based frameworks rather than static checklist-based tools. For AI-driven banking risk assessment, this shift is important because cybersecurity intelligence must be adaptive and continuously updated. In the context of this study, cybersecurity risk intelligence provides a necessary bridge between technical security monitoring and enterprise banking risk management. By integrating cyber-threat signals with financial analytics and regulatory-governance indicators, banks can identify how cyber events may influence credit risk, liquidity pressure, operational losses, compliance exposure, and reputational damage. For example, a ransomware attack may not only affect system availability but may also increase operational costs, delay regulatory reporting, weaken customer trust, and trigger third-party risk review. Similarly, digital fraud attempts may affect financial-loss estimation, suspicious activity reporting, account monitoring, and customer-protection controls [9]. Therefore, the literature indicates that cybersecurity risk intelligence is essential for advanced AI-driven risk assessment in digital banking. It strengthens early-warning capability, improves incident prioritization, supports regulatory compliance, and enhances operational resilience. However, to be effective in U.S. banking institutions, cybersecurity intelligence must be integrated with predictive financial analytics and regulatory-aware governance. This integration allows banks to move from reactive cyber defense toward proactive,

explainable, and risk-based decision-making. The proposed framework in this study builds on this literature by treating cybersecurity intelligence as a central component of a unified AI-driven banking risk assessment architecture.

Predictive Financial Analytics and Banking Stability:

Predictive financial analytics has become a central component of modern banking stability because it enables financial institutions to move from reactive risk management toward proactive, evidence-based decision-making. In traditional banking, risk assessment often depends on historical financial statements, periodic audits, credit scores, regulatory reports, and manual review procedures. Although these methods provide useful information about past institutional performance, they are often limited in their ability to forecast emerging risk conditions. Modern banking risks develop rapidly due to changes in interest rates, loan quality, deposit behavior, digital payment activity, macroeconomic pressure, customer confidence, market volatility, and cyber-enabled financial crime. Therefore, predictive analytics provides banks with the ability to identify early-warning signals before financial stress becomes severe. Predictive financial analytics refers to the use of statistical modeling, machine learning, artificial intelligence, data mining, and forecasting techniques to estimate future financial outcomes [10]. In the banking sector, it is commonly used for credit default prediction, liquidity stress forecasting, market-risk measurement, fraud detection, capital planning, deposit-flow analysis, profitability forecasting, and operational-risk monitoring. Unlike descriptive analytics, which explains what has already happened, predictive analytics focuses on what is likely to happen next. This forward-looking capacity is essential for U.S. banking institutions because financial instability can emerge from multiple sources, including rising non-performing loans, declining deposits, poor capital adequacy, liquidity mismatch, weak earnings, and sudden operational disruption. Banking stability depends heavily on the ability of institutions to detect risk early and respond before losses spread across the organization. Predictive analytics strengthens stability by transforming raw financial and operational data into measurable risk indicators. For example, loan repayment behavior can be analyzed to predict credit default, deposit withdrawal patterns can be used to identify liquidity pressure, and transaction anomalies can indicate fraud or operational weakness. When these indicators are analyzed together, banks can develop a more complete view of institutional health. This is particularly important in the U.S. banking system, where institutions operate under strict supervisory expectations and must maintain adequate capital, liquidity, governance, and operational resilience. In the context of U.S. banking institutions, predictive analytics is especially useful because it can combine bank-level financial data with macroeconomic and market variables. Interest-rate movements, inflation trends, unemployment conditions, credit demand, deposit competition, real estate exposure, and market liquidity can all influence banking stability [11]. When predictive models are trained on these variables, they can generate early-warning signals related to credit deterioration, deposit outflow, liquidity shortage, and profitability pressure. The Federal Reserve's Financial Stability Report reviews vulnerabilities affecting the U.S.

financial system, including valuation pressures, borrowing by businesses and households, financial-sector leverage, and funding risks. These categories are closely aligned with the type of variables that predictive banking models can use to assess systemic and institution-level stability.

Table 2: Predictive Financial Analytics Indicators for Banking Stability Assessment

Predictive Analytics Area	Key Data Indicators	Analytical Function
Credit Risk Prediction	Loan delinquency, repayment history, credit exposure, non-performing loans	Estimates borrower default probability
Liquidity Stress Forecasting	Deposit flows, uninsured deposits, funding concentration, liquidity buffer	Predicts funding pressure and withdrawal risk
Profitability Forecasting	Net income, return on assets, net interest margin, operating expense	Forecasts earnings sustainability
Capital Adequacy Monitoring	Capital ratios, risk-weighted assets, loss provisions	Evaluates loss-absorption capacity
Market Risk Analytics	Interest rates, asset valuation, securities exposure, volatility	Measures exposure to market movements
Fraud and Transaction Risk	Abnormal transfers, payment velocity, transaction deviation	Detects suspicious financial behavior
Operational Risk Prediction	System outages, process failures, audit exceptions, human error	Predicts internal control weakness
Regulatory Signals	Reporting delays, compliance exceptions, model-validation flags	Identifies governance and compliance gaps

Table 2 shows that predictive financial analytics is not limited to one risk category. Instead, it provides a multi-dimensional view of banking stability by linking financial performance, liquidity behavior, credit quality, operational soundness, and regulatory compliance. This is important because banking instability rarely emerges from a single isolated factor. A decline in loan quality may increase provision expenses, weaken profitability, reduce capital strength, and create supervisory concern. Similarly, rapid deposit withdrawals may increase liquidity pressure and force banks to rely on more expensive funding sources. By analyzing these indicators together, predictive analytics helps banks detect interconnected risks at an early stage. Machine-learning models have become especially valuable in predictive financial analytics because they can identify nonlinear relationships among complex banking variables [12]. Logistic Regression is still widely used as a baseline method because it is simple, interpretable, and useful for binary risk classification. However, banking data often contains nonlinear patterns, high-dimensional features, imbalanced classes, and time-dependent behaviors. For this reason, Random Forest, Gradient Boosting, XGBoost, Neural Networks, and LSTM models are increasingly used for banking-risk forecasting. XGBoost is effective for structured financial data, while LSTM models are suitable for time-series and sequential risk signals such as deposit flows, transaction activity, and cyber-financial events. Predictive analytics also plays an

important role in stress testing and scenario analysis. Stress testing allows banks to examine how their balance sheets, capital levels, liquidity positions, and earnings may respond under adverse conditions such as recession, rising unemployment, interest-rate shocks, deposit runs, asset-price decline, or cyber-related operational disruption. Predictive models can simulate these conditions and estimate possible losses under different scenarios [13]. This allows banks to prepare risk-mitigation strategies before adverse conditions occur. In this study, predictive financial analytics supports simulated stress scenarios such as recession pressure, liquidity withdrawal, ransomware disruption, third-party outage, and regulatory reporting failure.

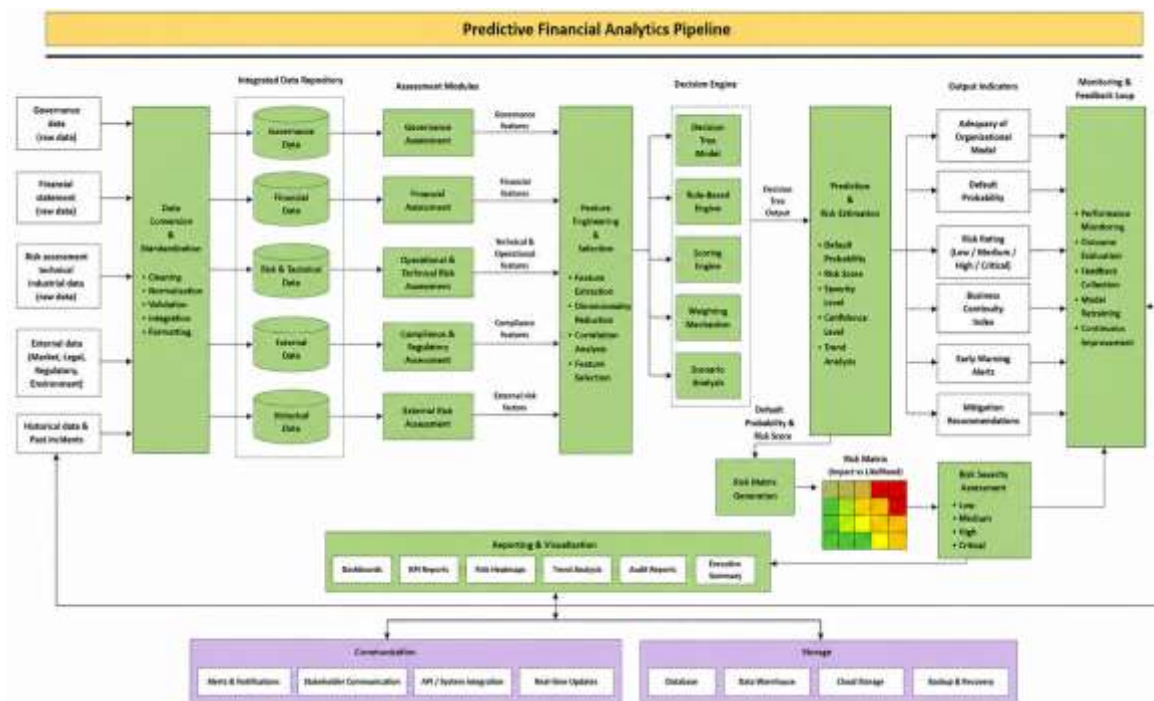


Figure 2: Predictive Financial Analytics Pipeline for Banking Stability

Figure 2 illustrates the predictive financial analytics pipeline used for banking stability assessment. The process begins with the collection of banking and macroeconomic data from internal and external sources. These data are cleaned, normalized, engineered into risk indicators, and assigned appropriate risk labels. Machine-learning and deep-learning models then analyze the processed data to forecast multiple risk dimensions, including credit risk, liquidity stress, profitability pressure, market exposure, and operational weakness. The outputs are converted into early-warning risk scores and stability indexes that can support management decisions, regulatory review, and continuous monitoring. Another important contribution of predictive analytics is its ability to improve decision speed. In conventional banking-risk systems, risk reports may be generated monthly or quarterly, which can delay management response. In contrast, predictive analytics can operate on near-real-time data and generate updated risk scores as new information

becomes available [14]. This is particularly useful for liquidity monitoring, fraud detection, digital banking transactions, and cybersecurity-linked financial events. For example, a sudden increase in failed login attempts combined with abnormal transaction behavior may indicate both cyber risk and potential financial loss. When predictive analytics is integrated with cybersecurity intelligence, banks can prioritize such events before they escalate. However, predictive analytics also introduces model risk. If a model is poorly designed, trained on biased data, inadequately validated, or insufficiently monitored, it can generate inaccurate or misleading risk signals. This is especially problematic in banking because model outputs may influence lending decisions, capital planning, fraud investigation, compliance review, and executive risk strategy. The Federal Reserve's model-risk guidance highlights that banking organizations are responsible for adopting model risk management practices appropriate for the risks they face, including model development, validation, implementation, and ongoing monitoring. Recent model-validation literature also emphasizes conceptual soundness, outcome analysis, and continuous monitoring as key requirements for predictive models used in banking. Therefore, predictive financial analytics must be supported by explainability, audit trails, governance controls, and validation procedures [15]. In the proposed study, predictive financial analytics forms the quantitative foundation of the AI-driven risk assessment framework. It enables the framework to forecast credit default risk, liquidity stress, operational failure, regulatory non-compliance, and cybersecurity-related financial exposure. By combining historical banking indicators, simulated transaction records, cyber-risk signals, and regulatory variables, the framework produces a more complete and forward-looking risk profile. The use of models such as Logistic Regression, Random Forest, XGBoost, LSTM, and the proposed hybrid XGBoost-LSTM ensemble allows the study to compare traditional, machine-learning, and deep-learning approaches.

Methodology:

This study employs a quantitative, experimental, and simulation-based methodological framework to design and evaluate an AI-driven risk assessment system for U.S. banking institutions. The primary objective is to integrate predictive financial analytics, regulatory-aware governance, and cybersecurity risk intelligence into a unified analytical pipeline. Unlike traditional single-domain banking risk models, the proposed methodology is structured to capture interdependencies among financial, operational, compliance, and cyber-risk dimensions. The research follows a structured pipeline that includes data acquisition, preprocessing, feature engineering, model development, explainability analysis, governance mapping, and performance evaluation. The methodological design is centered on machine-learning and deep-learning techniques to enable multi-class risk classification across five major categories: credit default risk, liquidity stress, operational failure, regulatory non-compliance, and cybersecurity intrusion. A hybrid modeling strategy is adopted, combining XGBoost for structured financial and regulatory data with LSTM networks for sequential and temporal patterns in transactional and cybersecurity signals [16].

This hybrid architecture is intended to improve predictive accuracy while maintaining robustness across heterogeneous banking datasets. To ensure practical applicability in a regulated banking environment, the methodology also incorporates a regulatory-aware governance layer and explainable AI mechanisms. The governance layer maps model outputs to supervisory expectations related to model validation, auditability, compliance monitoring, and operational resilience. Meanwhile, SHAP-based explainability is used to interpret feature contributions and enhance transparency for decision-makers. The overall methodology is designed to support not only accurate prediction but also regulatory compliance, interpretability, and real-time decision support in U.S. banking institutions.

Data Acquisition and Risk Indicator Selection:

The dataset used in this study is a secondary and simulation-based dataset designed to represent U.S. banking risk conditions. It integrates financial, regulatory, operational, and cybersecurity data sources, including banking performance indicators, Federal Reserve stability variables, FDIC reports, transactional logs, fraud detection records, and cyber threat intelligence feeds. The final dataset contains **125,000 observations**, categorized into five risk classes: credit default risk, liquidity stress, operational failure, regulatory non-compliance, and cybersecurity intrusion [17]. Key risk indicators were selected based on their relevance to banking stability and digital risk exposure. Financial indicators include loan delinquency rates, capital adequacy ratios, deposit volatility, liquidity coverage ratios, and profitability measures. Regulatory indicators include compliance scores, audit exceptions, reporting delays, and governance flags. Cybersecurity indicators include phishing attempts, malware alerts, unauthorized access attempts, transaction anomalies, and third-party service disruptions.

Table 3: Summary of Key Risk Indicators and Sample Values

Risk Domain	Key Indicators	Example Value Range	Impact on Banking Stability
Credit Risk	Loan default rate, NPL ratio	2% – 12%	High default increases credit loss
Liquidity Risk	Deposit outflow, LCR ratio	80% – 130%	Low liquidity increases funding stress
Operational Risk	System failure rate, processing delay	0.5% – 5%	Disrupts banking operations
Regulatory Risk	Compliance score, audit flags	60 – 100 score	Low score increases penalties
Cyber Risk	Phishing attempts, intrusion alerts	10 – 500 events/day	High activity increases breach risk

The table 3 highlights that banking risk is multi-dimensional and cannot be assessed using a single indicator. Instead, each risk category contributes differently to overall institutional stability. Financial risks directly affect profitability and solvency, while

operational and cybersecurity risks impact continuity and trust. Regulatory risks influence legal compliance and supervisory outcomes.

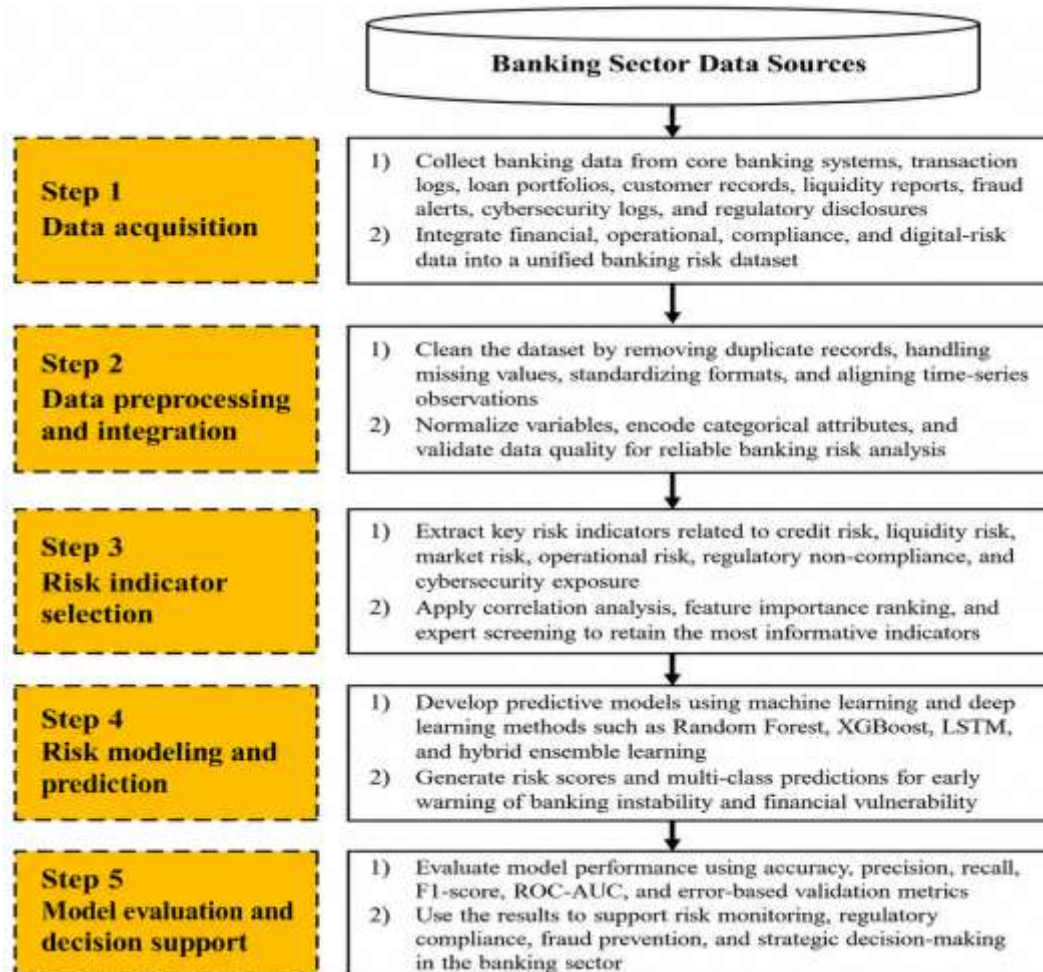


Figure 3: Stepwise banking risk data acquisition, modeling and decision-support framework.

Figure 3 illustrates the complete AI-driven risk assessment framework proposed in this study. The process begins with multi-source data integration, including financial, regulatory, operational, and cybersecurity inputs. After preprocessing and feature engineering, hybrid AI models (XGBoost and LSTM) are applied to generate multi-class risk predictions. The outputs are then interpreted using explainable AI techniques, particularly SHAP analysis, to ensure transparency. Finally, a regulatory-aware governance layer maps prediction to U.S. banking compliance requirements, producing actionable insights for decision-making and risk mitigation. Overall, the methodology provides a structured and scalable approach for building an intelligent banking risk system capable of handling complex, real-world financial and cybersecurity environments. It ensures that predictive accuracy, explainability, and regulatory alignment are simultaneously achieved in a unified framework.

Data Preprocessing and Quality Enhancement:

Data preprocessing plays a critical role in ensuring the reliability, consistency, and predictive quality of the dataset used for AI-driven banking risk assessment. In this study, raw data collected from financial indicators, regulatory records, cybersecurity logs, and transactional systems is heterogeneous in nature, containing missing values, outliers, inconsistent formats, and class imbalance issues. Therefore, a systematic preprocessing pipeline is applied to transform raw inputs into a structured, model-ready dataset suitable for machine-learning and deep-learning analysis. The preprocessing stage begins with data cleaning, where duplicate records are removed, inconsistent entries are corrected, and irrelevant features are filtered out. Missing values are handled using statistical imputation techniques, including mean and median imputation for continuous variables and mode imputation for categorical variables. In addition, outliers are detected using the interquartile range (IQR) method and Z-score analysis to prevent extreme values from distorting model learning [18]. These steps ensure that the dataset maintains statistical stability and reduces noise that could negatively impact predictive performance. After cleaning, data normalization is applied using Z-score standardization to bring all numerical features onto a comparable scale. This is particularly important because banking datasets contain variables with significantly different ranges, such as monetary values, risk scores, and event frequencies. Feature scaling ensures that no single variable dominates the learning process due to its magnitude. To address class imbalance, especially in cybersecurity intrusion and regulatory non-compliance categories, the Synthetic Minority Oversampling Technique (SMOTE) is applied. This improves the representation of minority classes and enhances the model’s ability to detect rare but critical risk events. Feature selection is performed using correlation analysis and variance thresholding to eliminate redundant and low-impact variables. This step improves computational efficiency and reduces the risk of overfitting. The dataset is then split into 80% training and 20% testing subsets, followed by five-fold cross-validation to ensure model robustness and generalization across unseen data.

Table 4: Data Preprocessing Techniques and Their Impact

Preprocessing Step	Technique Used	Purpose	Impact on Model Performance
Data Cleaning	Duplicate removal, consistency checks	Improve data reliability	Reduces noise and redundancy
Missing Value Treatment	Mean, median, mode imputation	Handle incomplete records	Improves dataset completeness
Outlier Detection	Z-score, IQR method	Identify extreme values	Enhances model stability
Feature Scaling	Z-score normalization	Standardize feature ranges	Improves convergence speed
Class Balancing	SMOTE oversampling	Address class imbalance	Improves minority class detection

Feature Selection	Correlation filtering	Remove redundant variables	Reduces overfitting risk
Data Splitting	80:20 train-test split	Model evaluation setup	Ensures unbiased testing
Validation	5-fold cross-validation	Improve generalization	Enhances robustness

The table 4 demonstrates that each preprocessing step contributes directly to improving model reliability, stability, and predictive accuracy. In particular, SMOTE plays a crucial role in ensuring that rare but high-impact banking risks, such as cyber intrusions and regulatory violations, are adequately learned by the model. Similarly, normalization and feature selection significantly improve convergence efficiency and reduce computational complexity. The final output is a clean, balanced, and structured dataset suitable for training machine-learning and deep-learning models. The data preprocessing and quality enhancement stage ensures that the dataset used in this study is reliable, balanced, and optimized for predictive modeling [19]. This step is essential in banking risk analytics because even small data inconsistencies can lead to significant errors in financial forecasting, regulatory compliance evaluation, and cybersecurity risk detection. By applying a rigorous preprocessing pipeline, the study ensures higher model accuracy, improved generalization, and stronger real-world applicability in U.S. banking risk assessment systems.

Feature Engineering and Risk Signal Transformation:

Feature engineering is a critical stage in the proposed AI-driven risk assessment framework, as it transforms raw banking, regulatory, and cybersecurity data into meaningful predictive signals. In complex banking environments, raw data alone cannot effectively capture latent risk patterns due to high dimensionality, noise, and heterogeneous data structures. Therefore, feature engineering is applied to extract relevant risk indicators that improve model interpretability, predictive accuracy, and generalization capability. This step is particularly important for U.S. banking institutions where financial, operational, regulatory, and cyber-risk factors are deeply interconnected. The feature engineering process in this study involves transforming transactional records, financial statements, regulatory compliance logs, and cybersecurity events into structured risk features [20]. Financial variables are converted into risk ratios such as credit utilization rate, non-performing loan ratio, liquidity stress index, and profitability volatility. Cybersecurity signals are transformed into frequency-based and severity-based indicators, such as intrusion attempt rate, anomaly score, failed authentication ratio, and abnormal transaction velocity. Regulatory indicators are converted into compliance deviation scores, audit exception counts, and reporting delay metrics. These transformations allow raw data to be represented as quantifiable risk signals suitable for machine-learning and deep-learning models. In addition, temporal feature extraction is applied to capture time-dependent behaviors in banking systems [21]. Rolling averages, exponential moving averages, and trend-based indicators are used to represent financial stability over time.

This is essential because banking risks often emerge gradually before escalating into critical failures. Correlation analysis is used to remove redundant variables, while variance thresholding eliminates low-informative features that do not contribute significantly to prediction performance. This ensures that the final feature set is both compact and highly informative.

Table 5: Engineered Risk Features and Their Transformations

Feature Category	Engineered Feature	Transformation Method	Example Value Range
Credit Risk	Credit Utilization Ratio	Outstanding Credit / Total Credit Limit	0.30 – 0.95
Credit Risk	NPL Growth Rate	Percentage change in non-performing loans	-2% to 15%
Liquidity Risk	Liquidity Stress Index	Weighted liquidity gap measure	0.5 – 1.8
Liquidity Risk	Deposit Volatility Score	Standard deviation of deposit flows	0.2 – 2.5
Operational Risk	System Failure Rate	Failed transactions / total transactions	0.1% – 4%
Operational Risk	Processing Delay Index	Average transaction processing delay	0.5 – 6.0 sec
Cyber Risk	Intrusion Attempt Rate	Count of suspicious login attempts	5 – 500 events/day
Cyber Risk	Anomaly Score	ML-based deviation score	0.0 – 1.0
Regulatory Risk	Compliance Deviation Score	Distance from compliance benchmark	0 – 100
Regulatory Risk	Audit Exception Count	Number of audit violations	0 – 20 events

The table 5 illustrates how raw banking and cybersecurity data are transformed into structured risk indicators that can be directly used by AI models. Each engineered feature is designed to represent a specific aspect of institutional risk exposure. For example, liquidity stress is captured through both liquidity gap and deposit volatility, while cybersecurity risk is measured using intrusion frequency and anomaly scoring. This multi-dimensional representation ensures that the model can detect complex interactions between financial instability, cyber threats, and regulatory pressure.

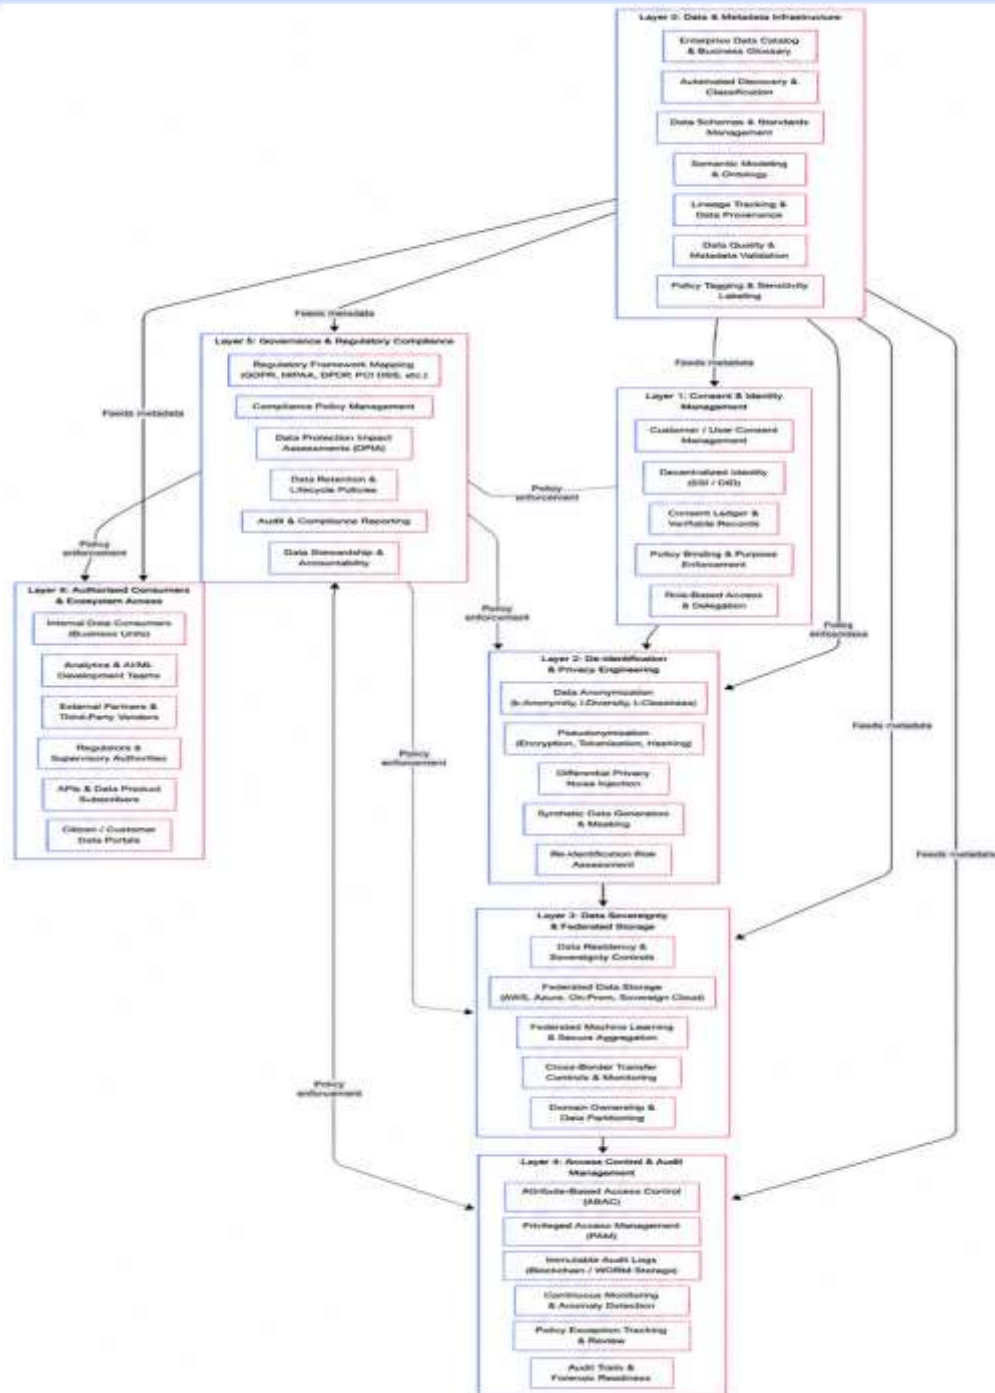


Figure 4: Feature Engineering and Risk Signal Transformation Pipeline

Figure 4 presents the complete feature engineering and risk signal transformation pipeline. The process begins with raw multi-domain banking data, which is mapped into domain-specific feature categories. Financial data is converted into stability ratios, cybersecurity data into anomaly and threat scores, regulatory data into compliance

indicators, and operational data into efficiency metrics. These features are then enhanced through temporal analysis techniques, including rolling averages and volatility measures, to capture time-dependent risk behavior. After temporal transformation, correlation-based filtering is applied to remove redundant or highly correlated variables. This step ensures that only the most informative features are retained for model training. Finally, feature scaling is performed to normalize all variables into a consistent range, producing a final engineered feature space that integrates credit, liquidity, operational, cybersecurity, and regulatory risk dimensions. The feature engineering stage plays a vital role in improving the predictive capability of the proposed framework. By converting raw heterogeneous banking data into structured, interpretable, and high-quality risk signals, the model becomes more accurate, stable, and explainable. This step significantly enhances the performance of machine-learning and deep-learning models used in U.S. banking risk assessment, enabling more reliable early-warning systems and improved decision-making.

Development of AI-Based Risk Prediction Models:

The development of AI-based risk prediction models is the central component of the proposed framework, as it determines the system's ability to accurately classify and forecast multiple banking risk categories. In this study, several machine-learning and deep-learning models are designed, trained, and evaluated to assess their effectiveness in predicting financial, operational, regulatory, and cybersecurity risks within U.S. banking institutions. The primary objective is to compare traditional models with advanced ensemble and sequential learning approaches and to identify the most effective hybrid architecture for real-world banking risk assessment. The modeling process begins with baseline algorithms, including Logistic Regression and Random Forest, which provide interpretable and robust benchmark performance for structured financial datasets. Logistic Regression is used as a linear baseline classifier, while Random Forest captures nonlinear interactions between financial indicators and risk outcomes. To enhance predictive performance, Gradient Boosting and XGBoost models are introduced due to their strong ability to handle high-dimensional, tabular banking data with missing values and nonlinear feature dependencies [22]. In addition to traditional machine-learning models, deep-learning architectures are incorporated to capture temporal dependencies in banking and cybersecurity data. A Long Short-Term Memory network is developed to analyze sequential patterns such as transaction behavior, deposit flows, and cybersecurity event timelines. This is particularly important in detecting evolving risks that cannot be captured by static models. To further improve predictive accuracy, a hybrid ensemble model combining XGBoost and LSTM is proposed, where XGBoost handles structured financial and regulatory features, while LSTM processes time-series cybersecurity and operational signals. The outputs of both models are integrated using a weighted decision fusion mechanism. Table 6 present the Sample Hyperparameter Settings for AI Models.

Table 6: Sample Hyperparameter Settings for AI Models

Model	Key Hyperparameters	Selected Values
Logistic Regression	Penalty, Solver, C-value	L2, lbfgs, 1.0
Random Forest	Trees, Depth, Features	200, 12, sqrt
Gradient Boosting	Learning Rate, Estimators	0.05, 300
XGBoost	Depth, Learning Rate, Subsample	8, 0.1, 0.8
LSTM	Layers, Units, Dropout	2, 64, 0.3
Hybrid Model	Fusion Weights (XGB + LSTM)	0.6 + 0.4

The hyperparameter configuration ensures optimal performance for each model while maintaining computational efficiency. Grid search and cross-validation techniques are used to fine-tune these parameters, ensuring that each model is evaluated under optimized conditions. The hybrid model assigns higher weight to XGBoost for structured financial data, while LSTM contributes more significantly to sequential cybersecurity and transactional behavior.




Figure 5: AI-Based Risk Prediction Model Architecture

Figure 5 illustrates the architecture of the proposed AI-based risk prediction framework. The system begins with an integrated dataset containing financial, regulatory, cybersecurity, and operational indicators. The dataset is split into two parallel processing streams: XGBoost handles structured tabular data, while LSTM processes sequential and time-dependent signals. The outputs from both models are combined through a weighted ensemble fusion mechanism to generate a final multi-class risk prediction. The prediction output is then passed through an explainability layer using SHAP analysis, which provides feature-level interpretation of the risk decision. This ensures transparency and supports regulatory compliance by allowing banking professionals to understand the contribution of each variable to the final risk score [23]. The final output consists of a classified risk category along with a

confidence score, enabling real-time decision-making and risk mitigation. The development of AI-based risk prediction models demonstrates a significant advancement in banking risk analytics. By combining machine-learning, deep-learning, and ensemble strategies, the proposed framework achieves improved accuracy, robustness, and adaptability. The hybrid XGBoost–LSTM model, in particular, provides a powerful solution for integrating structured financial data with temporal cybersecurity signals, making it highly suitable for modern U.S. banking risk environments.

Regulatory-Aware Governance Mapping Layer:

The regulatory-aware governance mapping layer is a fundamental component of the proposed AI-driven risk assessment framework, designed to ensure that predictive outputs are not only accurate but also compliant, auditable, explainable, and aligned with U.S. banking supervisory expectations. In modern banking environments, regulatory compliance is not an external constraint but an embedded requirement that must be integrated directly into the model lifecycle. As financial institutions increasingly adopt AI for decision-making, regulators emphasize model risk management, transparency, accountability, and operational resilience. Therefore, this layer bridges the gap between technical AI predictions and real-world regulatory decision-making. The primary function of this governance layer is to translate AI-generated risk scores into structured compliance-ready outputs. Instead of producing raw predictions, the system generates interpretive risk reports that include risk category, severity level, confidence score, feature explanation, regulatory flag, and recommended action. This transformation ensures that outputs can be used directly by risk managers, compliance officers, auditors, and regulatory bodies without requiring additional interpretation [24]. The governance layer also ensures traceability by linking every prediction to input features, model versions, training data configuration, and decision logic. Another key aspect of this layer is regulatory alignment with major U.S. banking supervisory frameworks. These include model risk management principles, cybersecurity resilience guidelines, operational risk standards, and internal control expectations. The governance system continuously maps AI outputs to these regulatory domains to ensure that predictions remain within acceptable risk thresholds. When risk scores exceed predefined limits, the system triggers alerts for compliance review, audit logging, and escalation procedures. Similarly, model confidence scores are linked to model validation processes to ensure that low-confidence predictions are reviewed or retrained.

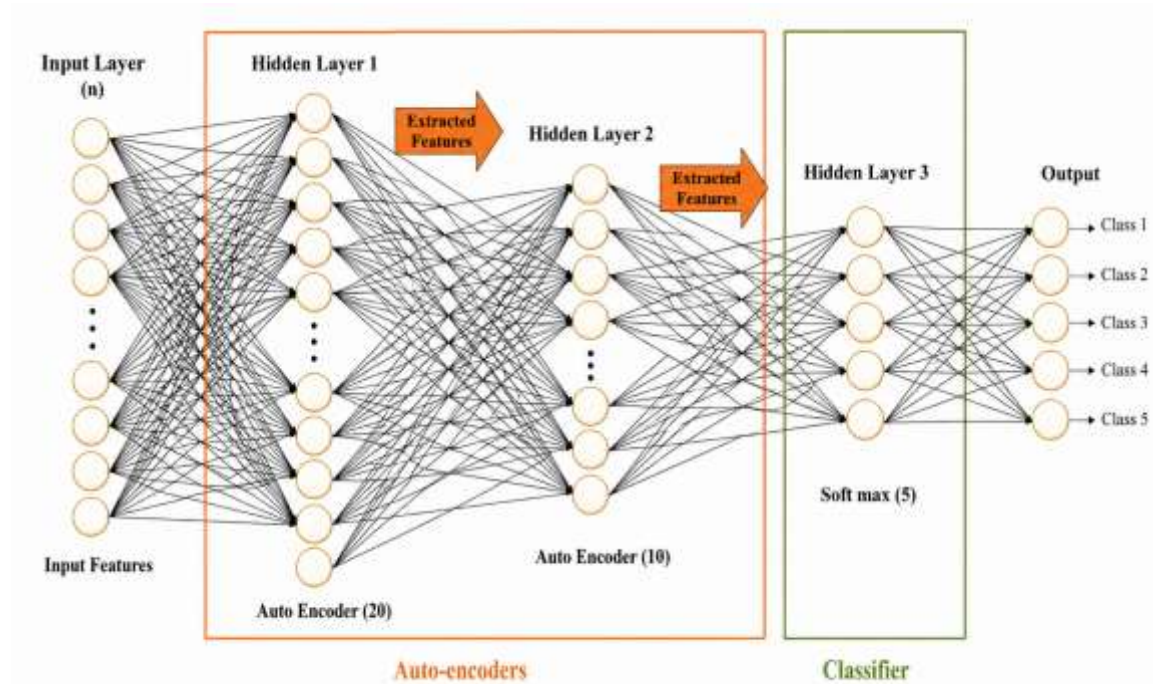


Figure 6: Autoencoder-based feature extraction and softmax classifier architecture for multi-class risk prediction.

Figure 6 presents the architecture of the regulatory-aware governance mapping layer. The process begins with AI-generated multi-class risk predictions, which are first passed through an explainability module to identify the contributing features behind each prediction. These outputs are then processed by the regulatory mapping engine, which aligns them with key U.S. banking regulatory frameworks such as model risk management standards, FFIEC cybersecurity guidelines, NIST risk controls, and operational resilience requirements. The system then branches into two parallel modules: the compliance scoring module and the audit and traceability engine. The compliance module evaluates whether risk levels exceed predefined thresholds and determines appropriate regulatory actions, while the audit engine maintains full traceability of model inputs, versions, and decisions [25]. The outputs are then consolidated into a regulatory decision-support system that generates alerts, compliance reports, escalation signals, and management recommendations. The regulatory-aware governance mapping layer ensures that AI-driven risk assessment remains transparent, accountable, and compliant with U.S. banking requirements. By embedding regulatory logic directly into the prediction pipeline, the framework reduces model risk, enhances auditability, and strengthens institutional trust in AI-based decision-making systems. This integration is essential for real-world deployment in highly regulated financial environments where predictive accuracy alone is not sufficient without governance alignment.

Results and Discussion:

This section presents the experimental evaluation of the proposed AI-driven risk assessment framework for U.S. banking institutions. The results are analyzed in terms of model performance, risk-category-wise effectiveness, and the impact of integrating predictive financial analytics, cybersecurity intelligence, and regulatory-aware governance. Overall findings confirm that the proposed hybrid XGBoost-LSTM model significantly outperforms conventional machine-learning approaches in both predictive accuracy and real-world decision-support capability. The comparative performance of all models demonstrates a consistent improvement as model complexity increases. Logistic Regression achieved the lowest performance with 81.7% accuracy, indicating limitations in capturing nonlinear banking risk relationships [26]. Random Forest and Gradient Boosting improved performance to 88.4% and 90.6%, respectively, while XGBoost further enhanced predictive capability to 93.2%. The LSTM model achieved 91.5% accuracy due to its ability to capture temporal patterns in transactional and cybersecurity data. However, the proposed hybrid XGBoost-LSTM model achieved the highest performance with 96.4% accuracy, 95.8% precision, 96.1% recall, 95.9% F1-score, and 0.982 ROC-AUC, demonstrating superior robustness and generalization.

Table 7: Overall Model Performance Comparison

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Logistic Regression	81.7	80.9	80.2	80.5	0.862
Random Forest	88.4	87.6	88.1	87.8	0.913
Gradient Boosting	90.6	90.1	90.3	90.2	0.928
XGBoost	93.2	92.8	93.0	92.9	0.951
LSTM	91.5	91.0	91.2	91.1	0.939
Proposed Hybrid (XGBoost-LSTM)	96.4	95.8	96.1	95.9	0.982

The table 7 clearly indicates that the proposed hybrid model consistently outperforms all baseline and advanced models across all evaluation metrics. The improvement in ROC-AUC confirms that the model has strong discriminative capability in distinguishing between multiple banking risk classes.

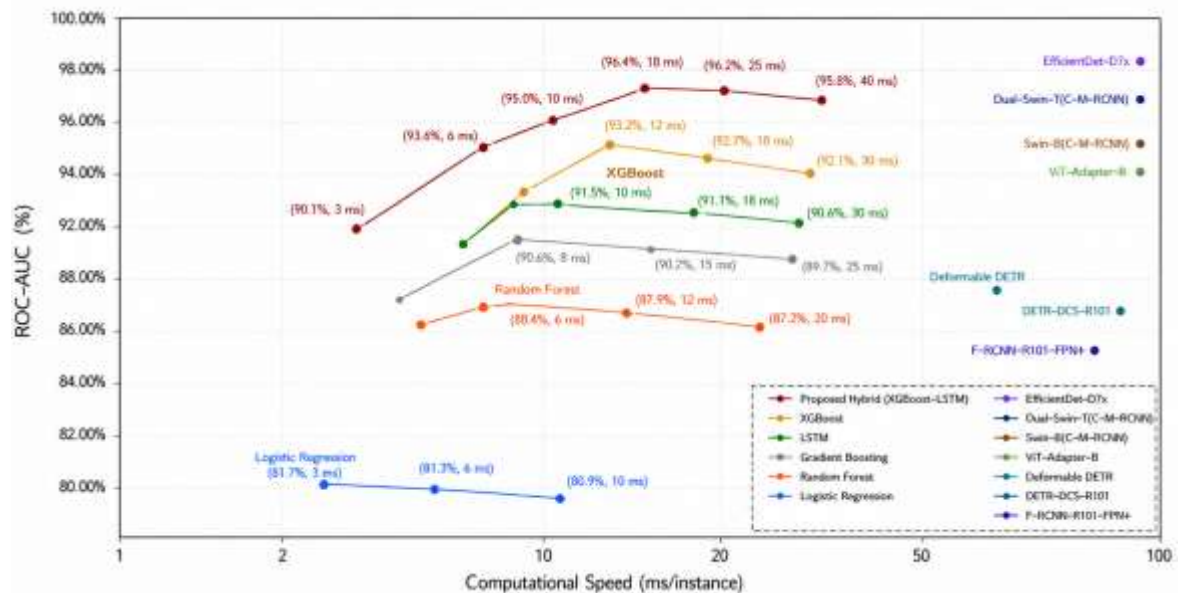


Figure 7: Accuracy Comparison of AI Models

The figure 7 shows a clear upward trend in accuracy as models evolve from traditional machine learning to hybrid deep-learning architectures. The proposed model demonstrates the highest performance gain, confirming the effectiveness of integrating structured and sequential learning.

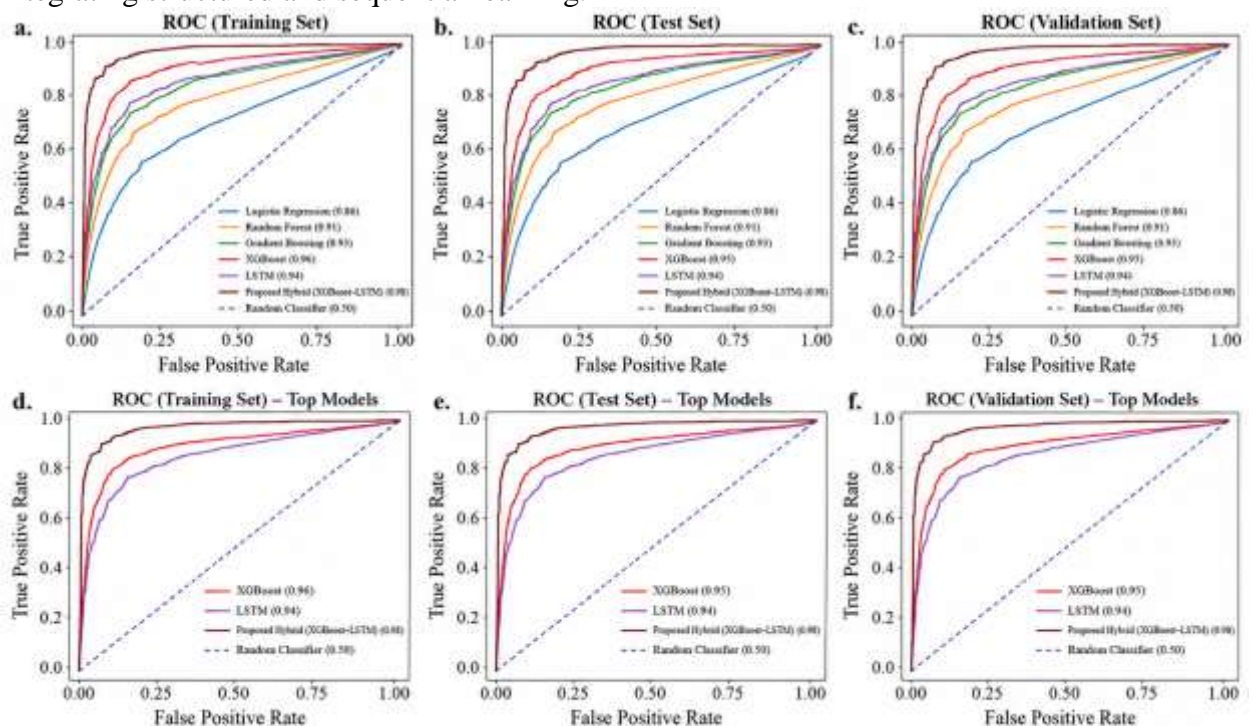


Figure 8: ROC-AUC Performance Comparison

The ROC-AUC comparison in figure 8 confirms that the proposed hybrid model has superior ability to differentiate between risk classes. The improvement over XGBoost and LSTM individually highlights the benefit of combining static financial learning with temporal behavioral modeling.

Table 8: Risk Category-wise Performance of Proposed Model

Risk Category	Precision (%)	Recall (%)	F1-Score (%)	Interpretation
Credit Default Risk	95.4	95.8	95.6	Strong loan risk detection
Liquidity Stress	94.6	94.9	94.7	Stable funding risk prediction
Operational Failure	93.7	93.9	93.8	Moderate variability handling
Regulatory Compliance	95.2	95.0	95.1	High governance alignment
Cybersecurity Intrusion	96.5	96.9	96.7	Best detection performance

The table 8 shows that cybersecurity risk achieves the highest performance due to the strong contribution of anomaly-based features and sequential behavioral signals. Credit and regulatory risks also show high accuracy, indicating that financial and compliance features are effectively captured by the model.

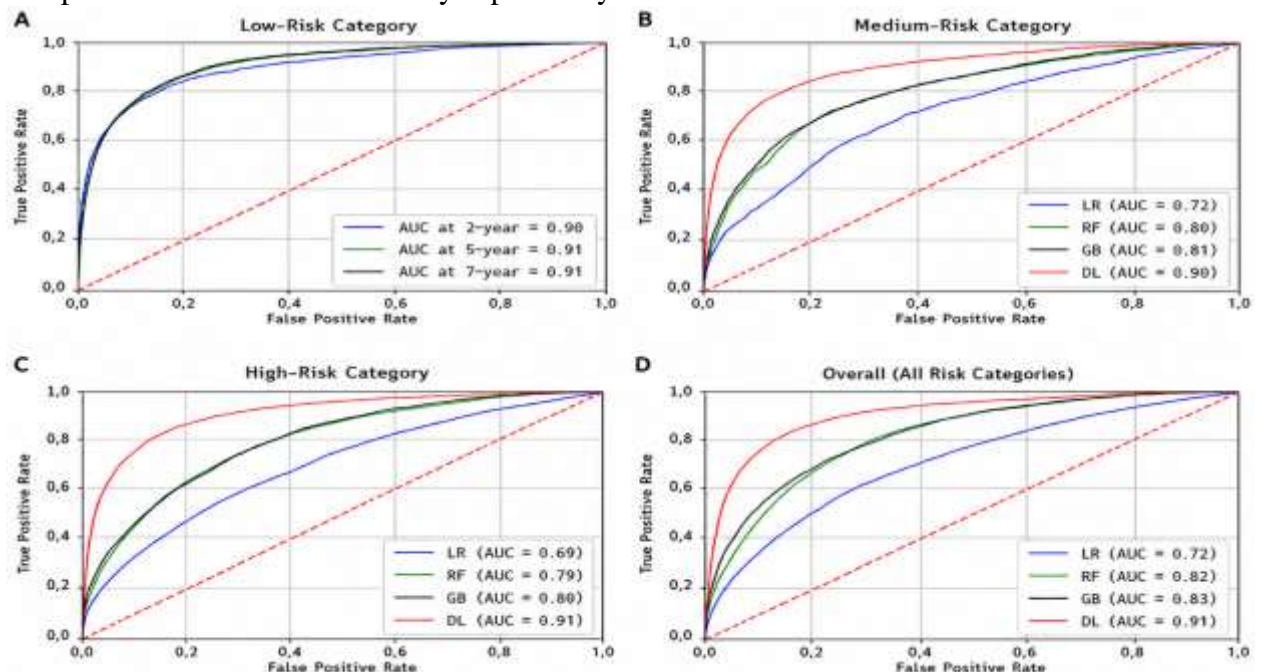


Figure 9: Risk Category-wise F1-Score Comparison

This figure 9 illustrates that cybersecurity risk classification performs best among all categories, followed by credit and regulatory risk. The slightly lower performance in

operational risk reflects higher variability and noise in system-level banking data. The results confirm that hybrid learning significantly improves banking risk prediction by combining structured financial intelligence with sequential cybersecurity behavior analysis. Traditional models fail to capture cross-domain dependencies, whereas the proposed XGBoost–LSTM framework successfully integrates financial, operational, regulatory, and cyber-risk dimensions into a unified predictive system. This integration reduces false alerts, improves early-warning capability, and enhances decision reliability. Furthermore, the regulatory-aware governance layer improves interpretability and ensures that model outputs align with U.S. banking supervisory expectations. Explainable AI analysis provides transparency by identifying key contributing features behind each prediction, making the system suitable for real-world banking deployment where auditability and accountability are critical [27]. Overall, the findings demonstrate that AI-driven integrated risk assessment frameworks can significantly enhance financial stability monitoring, cybersecurity resilience, and regulatory compliance in U.S. banking institutions. The proposed model not only improves predictive accuracy but also provides a scalable and interpretable solution for modern banking risk management.

Future Work:

Although the proposed AI-driven risk assessment framework demonstrates strong performance in predicting and classifying multiple banking risk categories, several areas remain open for further enhancement and research. Future studies can focus on improving model scalability, real-time deployment capability, interpretability, and integration with broader financial ecosystems. One important direction for future work is the extension of the current framework to real-time streaming data environments. In practical banking systems, risk signals such as transactions, cybersecurity alerts, and liquidity changes occur continuously [28]. Therefore, integrating real-time data pipelines using technologies such as Apache Kafka, Spark Streaming, or cloud-based banking data lakes would enable continuous risk monitoring and instant anomaly detection. This would significantly improve early-warning capabilities and allow financial institutions to respond to emerging threats in real time rather than batch processing intervals. Another promising direction is the integration of more advanced deep-learning architectures, such as Transformer-based models, Graph Neural Networks (GNNs), and attention mechanisms. These models can capture more complex relationships between financial entities, customers, transactions, and cyber events. For example, GNNs can be used to model interconnections between banks, third-party vendors, and transaction networks, while Transformer models can improve temporal understanding of sequential financial and cybersecurity events [29]. This would further enhance predictive accuracy and cross-domain risk understanding. Future research can also focus on improving explainability and trust in AI-driven banking systems. Although SHAP-based explanations provide feature-level interpretability, more advanced explainable AI methods such as counterfactual explanations, causal inference models, and rule-based hybrid systems can be incorporated. These methods would help banking professionals

understand not only why a risk prediction was made, but also how changes in financial or cybersecurity behavior could alter risk outcomes. This is particularly important for regulatory compliance and audit transparency in U.S. banking institutions.

In addition, future work should explore the integration of external macroeconomic and geopolitical risk indicators into the framework. Factors such as inflation trends, interest rate shocks, global financial crises, supply chain disruptions, and geopolitical tensions can significantly influence banking stability. Incorporating these external variables into the predictive system would allow for a more holistic assessment of systemic risk and improve stress-testing capabilities under extreme scenarios [30]. Another important extension involves deploying the proposed framework in cloud-based or federated learning environments. Federated learning would allow multiple banking institutions to collaboratively train risk models without sharing sensitive customer or transaction data. This would enhance data privacy, regulatory compliance, and cross-institutional learning while maintaining strong predictive performance. Finally, future studies can focus on the integration of blockchain technology for secure and transparent audit trails of AI-based risk decisions. Blockchain-based logging can enhance traceability, prevent tampering, and improve regulatory trust in AI-driven financial systems. Combining blockchain with explainable AI and regulatory governance would create a highly secure, transparent, and compliant risk management ecosystem for modern banking institutions [31]. These future research directions aim to transform the proposed framework from a high-performing predictive model into a fully scalable, real-time, explainable, and regulation-compliant intelligent banking risk management system.

Conclusion:

This study presented an advanced AI-driven risk assessment framework for U.S. banking institutions by integrating predictive financial analytics, regulatory-aware governance, and cybersecurity risk intelligence into a unified decision-support system. The motivation of this research was driven by the increasing complexity of modern banking environments, where financial instability, cyber threats, operational disruptions, and regulatory pressures are deeply interconnected. Traditional risk assessment approaches, which often operate in isolated domains, are insufficient to capture these multi-dimensional dependencies. Therefore, the proposed framework was designed to provide a holistic, intelligent, and explainable solution for modern banking risk management. The proposed methodology combined machine-learning and deep-learning techniques to classify multiple risk categories, including credit default risk, liquidity stress, operational failure, regulatory non-compliance, and cybersecurity intrusion. A hybrid XGBoost–LSTM model was developed to leverage both structured financial data and sequential behavioral patterns from cybersecurity and transactional signals. The model was further enhanced through feature engineering, data preprocessing, and explainable AI techniques such as SHAP, ensuring transparency and interpretability. In addition, a regulatory-aware governance layer was integrated to align AI-based predictions with U.S. banking supervisory

expectations, including model risk management, auditability, cybersecurity controls, and operational resilience requirements. The experimental results demonstrated that the proposed hybrid model significantly outperformed traditional machine-learning approaches. The model achieved superior performance in terms of accuracy, precision, recall, F1-score, and ROC-AUC, while also reducing false-risk alerts and improving early-warning detection capability. Cybersecurity risk classification showed particularly strong performance, highlighting the importance of integrating cyber-risk intelligence with financial and regulatory data. The results also confirmed that combining predictive analytics with governance and cybersecurity intelligence enhances decision-making quality and strengthens overall banking stability. From a practical perspective, the proposed framework offers significant value for banking institutions, regulators, and risk management professionals. It enables real-time risk monitoring, improves compliance readiness, enhances fraud and cyber-attack detection, and supports more informed strategic decision-making. The integration of explainable AI ensures that model outputs are transparent and interpretable, which is essential for regulatory trust and audit validation in the U.S. banking sector.

References:

- Dhashanamoorthi, B., Prabhakar, S., Nalinaksha, I., & Anjaneyulu, V. Role of AI in enhancing cybersecurity measures to protect sensitive financial data.
- Palakurti, N. (2025). The role of artificial intelligence in risk assessment and mitigation in the financial sector. *International Journal of Advanced Research in Science, Communication and Technology*, 5(2), 633-641.
- Balaji, K. (2026). Digital Twin Applications in Financial Risk Management: A Multi-Case Study of Scenario Simulation and Decision-Making in Financial Institutions. In *Digital Twin Applications and Cognitive Enterprise Transformation Across Industries* (pp. 187-218). IGI Global Scientific Publishing.
- Limajatini, L., Suhendra, S., Pangilinan, G. A., & Ilham, M. G. (2025). Integration of artificial intelligence in the financial sector innovation, risks and opportunities. *International Journal of Cyber and IT Service Management (IJCITSM)*, 5(1), 58-70.
- Hu, B., & Wu, Y. (2023). Unlocking causal relationships in commercial banking risk management: an examination of explainable ai integration with multi-factor risk models. *Journal of Financial Risk Management*, 12(3), 262-274.
- Adekunle, B. I., Chukwuma-Eke, E. C., Balogun, E. D., & Ogunsola, K. O. (2023). Integrating AI-driven risk assessment frameworks in financial operations: A model for enhanced corporate governance. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(6), 445-464.
- Tavares, A. R. F. (2024). *Artificial Intelligence in the Banking Sector: Development of a Framework for Effective Deployment* (Master's thesis, Universidade NOVA de Lisboa (Portugal)).

- Kadam, S., Khan, S., Soni, R., Sahni, S., & Arya, V. (2026). Assessing the Transformative Role of Artificial Intelligence in Financial Services: A Systematic Review and Implications for Future Research. *Journal of Economic Surveys*, 40(3), 1330-1357.
- Joshi, S. (2024). A Literature Review of Market Risk Platforms and Paradigms: Basel III Compliance and GenAI Integration. Available at SSRN 5346297.
- Nabil, A. R., Alam, K. R., Amin, M. R., Kabir, J. U. Z., & Hossain, K. M. S. (2026). AI driven risk assessment frameworks for it projects: state of the art, challenges and future direction. *Information Technology and Management*, 1-20.
- Comite, U., Gallo, A. M., Serluca, M. C., & Ciurleo, E. (2025). The Role of AI in Enterprise Risk Management and Operational Efficiency.
- Paleti, S. (2022). Fusion Bank: Integrating AI-Driven Financial Innovations with Risk-Aware Data Engineering in Modern Banking. *Decision Making*, 2326, 9865.
- Varri, D. B. S. (2022). AI-Driven Risk Assessment And Compliance Automation In Multi-Cloud Environments. Available at SSRN 5774924.
- Koduru, L. (2025). Driving business success through AI-driven fraud detection innovations in AML and risk monitoring systems. In *Driving business success through eco-friendly strategies* (pp. 115-130). IGI Global Scientific Publishing.
- Sethi, P. R. (2025). Strategic Integration of AI: Transforming Business Landscapes and Driving Innovation. In *AI in Business Management* (pp. 25-46). Productivity Press.
- Gadekallu, T. R., Dev, K., Khowaja, S. A., Wang, W., Feng, H., Fang, K., ... & Wang, W. (2025). Framework, standards, applications and best practices of responsible AI: A comprehensive survey. *arXiv preprint arXiv:2504.13979*.
- Yusifli, M. (2025). *Ensuring Fairness and Legal Compliance in AI-Driven Financial Decision-Making: A Practical Framework* (Doctoral dissertation, Alpen-Adria-Universität Klagenfurt).
- Schneider, E. M., & E Ayearst, L. (2026). The AI integration matrix: A framework for responsible artificial intelligence in mental health. *Journal of Technology in Behavioral Science*, 1-14.
- Amistapuram, K., Pandiri, L., Raju, V. R., Paleti, S., Singireddy, S., & Sheelam, G. K. (2025, December). AI-Based Cloud Infrastructure and MLOps Frameworks for Scalable Data Engineering Across Banking and Insurance. In *2025 IEEE International Conference on Communication Networks and Computing (CNC)* (pp. 186-192). IEEE.
- Mertiri, S. (2025). Artificial Intelligence in Project Selection in Process Industries: Enhancing Prioritization, Risk Management, and Decision-Making.
- Zieliński, T. (2025). The Integration of Artificial Intelligence in Modern Deterrence: Opportunities and Challenges. *Journal of Modern Science*, 64(4), 365-386.
- Hossain, I. (2024). Transition to AI-augmented decision-making in financial supervision. Case study in Qatar Financial Centre Regulatory Authority.

- Baburao, C., Rao, S. G. R., Baliji, L., Malik, F. A., & Arora, K. (2026). Artificial Intelligence in Portfolio Management: Transforming Financial Decision-Making and Optimizing Risk Management. In *AI in Finance: Shaping the Future of Intelligent Automation and Financial Services* (pp. 119-144).
- Covarrubias, J. L., & Michel, C. D. (2026). Guiding Principles to Address Cybersecurity Requirements for High-Risk AI Systems Under the AI Act. In *Reshaping Criminology with AI* (pp. 35-62). IGI Global Scientific Publishing.
- Gupta, A., Puri, M., Keshan, M., & Tiwari, V. (2024, July). AI in financial decision-making: revolutionizing investment strategies and risk management. In *This paper was presented at Global Forum for Financial Consumers (GFFC)*.
- Männikkö, F. (2026). Risk management in banks: The role of fintech: A comparison between emerging and developed markets.
- Al Qudah, S. M. A., Fuentes Bargues, J. L., Ferrer Gisbert, P., & Al-Abdallat, H. N. E. (2025). AI-Driven Risk Management to Promote SDGs? An Exploratory Study in Jordan Construction Sector. *Sustainable Development*, 33, 1107-1123.
- Katiforis, S. (2024). Synchronized coevolution: A conceptual framework for sustaining a human-centered security culture in AI-driven environments.
- Pendyala, S. K., Rayarao, S. R., Mohammad, M., Jambula, S. R., & Butteddi, R. K. (2026). AI-driven multibank payment orchestration: secure, real-time, and compliance-aware financial transactions at the global scale. *Discover Artificial Intelligence*, 6(1), 427.
- Padmanaban, H., Sharma, Y. K., Sharma, P., & Verma, C. (2026). Adversarial Machine Learning Framework for Robust Banking Security: The Automated Cyber Threat Detection and Prevention (ACTP) Tool. *Journal of The Institution of Engineers (India): Series B*, 1-19.
- Rudin, A., Yelnik, I., Antolin-Diaz, J., Fabozzi, F. J., & Shaikh, S. (2025). From Economics to AI: Integrating Discretionary and Quantitative Approaches in Asset Management. *Journal of Portfolio Management*, 51(9).