

A Composite AI Framework for Dynamic Operational Risk Management in Banking

Saeed Jafari Yazdanabad

M.A. in Economics, Islamic Banking Concentration, Department of Economics, Faculty of Economics,
University of Tehran, Tehran, Iran

Jafari.saeed@ut.ac.ir

Abstract

Operational risk management (ORM) is central to banking stability, particularly given the rapid growth of digital transactions, open banking, and stringent regulatory mandates like Basel III and IV. Traditional ORM relies on historic loss databases, periodic risk and control self-assessments (RCSA), and static key risk indicators (KRIs). These retrospective approaches struggle to identify complex process workarounds, internal collusion, and emergent systemic failures in real time. This paper proposes a comprehensive, five-layer Composite AI framework that integrates unsupervised anomaly detection (Isolation Forest and Deep Autoencoders), Natural Language Processing (BERT), Graph Neural Networks (GCN), and process mining. This architecture enables simultaneous monitoring of structured and unstructured data, employee-customer interaction networks, and transactional workflow deviations. The framework also incorporates an Explainable AI (XAI) layer using SHAP and LIME to address the "black-box" dilemma of deep learning models, smoothing the path for regulatory compliance. This paper demonstrates that the transition from passive sampling to a proactive, continuous, and explainable AI-driven ORM is a strategic imperative for resilient banking operations. The framework provides banking professionals with practical, actionable tools to modernize their risk management practices without requiring deep technical expertise in artificial intelligence.

Keywords: Operational Risk Management, Composite Artificial Intelligence, Process Mining, Anomaly Detection, Graph Neural Networks, Basel Accords

1. Introduction and Problem Statement

Modern banking environments are undergoing profound structural transformations. The global expansion of open banking, instant payment networks, microservice-based IT architectures, and decentralized platforms has exponentially scaled transaction volumes and operational speed (Arner et al., 2017). According to the official guidelines of the Basel Committee on Banking Supervision (BCBS), operational risk is defined as the risk of loss resulting from inadequate or failed internal processes, people, and systems, or from external events (Basel Committee on Banking Supervision, 2021). Unlike credit and market risks, which rely on structured quantitative datasets, market prices, and transparent default rates, operational risk is fundamentally decentralized, highly qualitative, unstructured, and deeply embedded within human workflows (Leo et al., 2019).

1.1. The Three Pillars of Traditional Operational Risk Management

To manage operational risk, contemporary financial institutions rely on three core pillars:

First, Loss Event Databases (LED) – These databases record historical operational loss incidents. However, they function purely as lagging indicators, registering events only after losses are finalized and reputational damage has occurred. By the time a loss is recorded, the opportunity for prevention has already passed.

Second, Risk and Control Self-Assessments (RCSA) – These are periodic assessments (typically semi-annual or annual) where business units evaluate their own risk exposures and control effectiveness. RCSAs are highly vulnerable to subjective human bias, memory lapses, and the reluctance of business units to transparently report control vulnerabilities within their domains. Staff may hesitate to report issues for fear of repercussions, and the infrequent nature of these assessments means risks can evolve substantially between review cycles.

Third, Key Risk Indicators (KRIs) – These are static, threshold-based metrics designed to signal when risk levels exceed acceptable boundaries. Traditional KRIs are built on rigid thresholds that generate high volumes of false alerts while failing to detect sophisticated, multi-stage internal fraud or structural process bypasses (Fritz-Morgenthal et al., 2022).

1.2. The Limitations of Traditional Approaches

These traditional methodologies suffer from severe systemic limitations that leave banks vulnerable to modern operational risks:

Retrospective Nature. Traditional tools only activate after an incident has occurred and losses have materialized. They offer no predictive capability and cannot prevent events from happening in the first place.

Reliance on Sampling. Traditional control systems rely heavily on random transactional sampling. This means the vast majority of transactions are never reviewed, allowing highly customized fraudulent schemes and policy infractions to pass through control networks unnoticed.

Inability to Detect Complex Patterns. Rule-based systems can only detect what they have been explicitly programmed to identify. They cannot recognize novel fraud patterns, emergent collusion networks, or sophisticated process workarounds that have never been seen before.

Underutilization of Unstructured Data. Approximately 80% of data relevant to operational risk—including internal audit logs, customer complaints, legal documents, and regulatory directives—remains unstructured and completely unutilized within classical ORM systems (Financial Stability Board, 2020; EBA, 2020).

Human Bias and Subjectivity. RCSA processes rely on the honesty, memory, and judgment of individual staff members. This introduces significant subjectivity and inconsistency into risk assessments.

1.3. The Dual Role of Artificial Intelligence

Recent research has revealed a fascinating paradox in the relationship between AI and operational risk. A study by the Federal Reserve Bank of Richmond and the Federal Reserve Bank of Boston demonstrates that banking organizations with higher AI investments are exposed to more operational risk (McLemore & Mihov, 2024). Using comprehensive supervisory data on operational losses from large U.S. bank holding companies combined with detailed company-level data on AI-skilled human capital, the study shows that banks with more AI investments suffer higher operational losses per dollar of total assets. The impact of AI investments on operational losses significantly varies by loss type and is driven by external fraud, client-related issues, and system failures. These losses stem not only from small, frequent incidents but also from severe, tail-risk events (McLemore & Mihov, 2024).

However, this does not mean banks should avoid AI. Rather, it highlights that **AI is a double-edged sword**: it simultaneously amplifies operational risks while offering powerful tools for enhanced risk mitigation (Deloitte & ORX, 2025). The risk-enhancing effect of AI is more pronounced for banks with weaker risk management practices (McLemore & Mihov, 2024). This suggests that the key differentiator is not whether a bank adopts AI, but **how** it adopts AI—with proper governance, controls, and oversight frameworks in place.

According to the ORX Perspectives on AI in Banking report, AI both amplifies and transforms well-known operational risks across information security (including cyber), technology and data, model risk, conduct, compliance, and third-party risk, while also introducing new "AI native" risks that are novel to the technology (Deloitte & ORX, 2025). At the same time, AI provides new opportunities to enhance operational risk processes and reduce risk exposures through advanced analytics and greater precision in process execution (Deloitte & ORX, 2025).

1.4. The Need for a New Approach

This paper addresses the gap between traditional risk management and the demands of modern digital banking by presenting a comprehensive, explainable Composite AI framework designed to transition banking institutions from retrospective operational risk mitigation to continuous, predictive, and dynamic risk management. The framework is designed to be understood and applied by banking professionals who may not have deep technical expertise in artificial intelligence, focusing on practical applications and tangible benefits rather than mathematical complexity.

2. How Artificial Intelligence Helps Banks Manage Operational Risk

To understand how AI can transform operational risk management, it is helpful to think of AI tools as specialized assistants, each with a unique capability. Just as a bank has different departments for different functions (credit, treasury, compliance), different AI techniques are suited to different types of operational risk challenges.

2.1. Unsupervised Learning: Finding the Unknown Unknowns

What it does: Unsupervised learning algorithms learn what "normal" looks like in banking operations and then flag anything that deviates from this norm. They do not need to be trained on examples of fraud or errors—they simply identify patterns that are unusual.

Why this matters for banks: The biggest challenge in operational risk is that severe loss events are rare. Banks may have very few examples of major fraud incidents or system failures to learn from. Unsupervised learning overcomes this by focusing on normal behavior rather than trying to predict rare events. When something unusual happens—even if it has never happened before—the system raises an alert.

The Isolation Forest technique works by isolating anomalies rather than profiling normal data. Imagine trying to find a needle in a haystack. Traditional approaches would examine every piece of hay to understand what "normal hay" looks like. Isolation Forest instead asks: how easily can this item be separated from everything else? Anomalies are easier to isolate because they are different (Liu et al., 2012). This makes Isolation Forest exceptionally effective for high-dimensional transaction data where anomalies may hide among millions of normal transactions.

Deep Autoencoders are neural networks trained exclusively on normal transaction patterns. They learn to compress normal transactions into a simplified representation and then reconstruct them. When a normal transaction is processed, the reconstruction is accurate. When an anomalous transaction is processed—one containing policy violations or unusual execution paths—the reconstruction has a high error rate, triggering an alert (Goodfellow et al., 2016).

Practical banking applications: A bank can deploy unsupervised learning to monitor all transactions continuously, flagging unusual patterns such as a sudden spike in transactions from a particular branch, unusual timing of activities, or atypical combinations of transaction types that might indicate fraud.

2.2. Natural Language Processing: Reading What People Write

What it does: Natural Language Processing (NLP) enables computers to read, understand, and extract meaning from human language. This includes emails, reports, complaints, regulatory documents, and audit findings.

Why this matters for banks: Approximately 80% of data relevant to operational risk is unstructured text—things people have written rather than numbers in spreadsheets (Financial Stability Board, 2020). Traditional risk management ignores this vast source of information. NLP allows banks to automatically process and extract insights from these texts.

The BERT (Bidirectional Encoder Representations from Transformers) model is a powerful NLP technique that understands the contextual relationships between words (Devlin et al., 2019). Unlike older methods that treat words in isolation, BERT understands that the meaning of a word depends on the words around it. This allows it to accurately interpret the meaning of complex documents.

Practical banking applications:

- **Automated loss event classification.** When an incident report is filed, BERT can automatically determine which of the seven Basel operational risk categories it belongs to, eliminating manual classification and ensuring consistency.

- **Regulatory compliance mapping.** When a new regulation is issued, BERT can compare it against the bank's internal policies and automatically identify gaps that need to be addressed.
- **Sentiment analysis of complaints.** By analyzing customer complaint texts, the system can identify emerging issues before they escalate into major problems.
- **Audit report analysis.** NLP can extract key findings from audit reports and flag recurring issues across different business units.

2.3. Graph Neural Networks: Connecting the Dots

What it does: Graph Neural Networks (GNNs) model relationships between entities—people, accounts, transactions, systems—as a network or "graph." They then analyze this network to find hidden connections and patterns.

Why this matters for banks: Operational risk events are rarely isolated. Fraud often involves collusion between multiple employees and external parties. System failures can cascade from one system to another. Traditional analysis treats each transaction or entity in isolation and misses these critical connections.

Graph Convolutional Networks (GCNs) are a specific type of GNN that can identify patterns in networks by analyzing how information flows between connected nodes (Kipf & Welling, 2017; Hamilton et al., 2017). If multiple entities are connected in unusual ways—for example, several new corporate accounts all being approved by the same loan officer using the same computer terminal—the GCN can detect this pattern even if each individual transaction appears normal.

Practical banking applications:

- **Collusion detection.** Identifying networks of employees and external parties working together to circumvent controls.
- **Contagion analysis.** Understanding how a failure in one system might cascade to others, identifying single points of failure.
- **Money laundering detection.** Following complex chains of transactions across multiple accounts and entities.
- **Vendor risk assessment.** Analyzing the network of third-party relationships to identify concentration risks.

2.4. Process Mining: Watching How Work Actually Gets Done

What it does: Process mining extracts data from system logs to reconstruct and analyze how business processes are actually executed, comparing the real process against the intended process.

Why this matters for banks: There is often a gap between how a process is supposed to work and how it actually works in practice. Employees may develop workarounds, skip steps, or find ways to bypass controls. Process mining reveals these deviations in real time.

How it works: Every transaction in a banking system generates an "event log" entry containing a Case ID (what is being processed), Activity Name (what action was taken), Timestamp (when it happened), and Resource ID (who did it) (van der Aalst, 2016). Process mining algorithms analyze these logs to reconstruct the actual sequence of steps in each process and compare it to the approved model.

Conformance checking calculates a "Fitness Score" for each process instance, measuring how closely the actual execution matches the intended design (de Medeiros et al., 2008). Any workflow step that violates established sequence boundaries—such as a financial transfer executed without the mandatory dual-signature approval—is immediately detected as a critical process bypass, generating a real-time operational risk alert.

Practical banking applications:

- **Control bypass detection.** Identifying when employees skip required approval steps.
- **Process optimization.** Finding bottlenecks and inefficiencies in operational workflows.
- **Fraud detection.** Identifying unusual patterns in how transactions are processed.
- **Compliance monitoring.** Ensuring that regulated processes are followed correctly.

3. Mapping AI Tools to Basel Operational Risk Categories

To demonstrate how these AI tools align with regulatory requirements, Table 1 maps the four core AI methodologies to the seven official Basel operational risk event types (Basel Committee on Banking Supervision, 2021).

Table 1: Alignment of Composite AI methodologies with Basel operational risk categories.

Basel Event Type	Key AI Tool/Model	What It Does	How It Helps the Bank
1. Internal Fraud	Isolation Forest + Graph Neural Networks	Analyzes access patterns and employee networks	Identifies unusual access times, unauthorized modifications, and employee collusion networks that might indicate internal fraud
2. External Fraud	Deep Autoencoders	Analyzes transaction patterns	Detects identity theft, account takeover attempts, and unusual payment flows in real time
3. Employment Practices & Safety	NLP (BERT)	Analyzes HR documents and complaints	Classifies human resource risks, identifies sentiment anomalies, and monitors workplace safety compliance
4. Clients, Products & Business Practices	Process Mining + NLP	Analyzes credit files, KYC documents, and advisory logs	Detects AML process gaps, mis-selling behaviors, and unauthorized advisory deviations
5. Damage to Physical Assets	Time-Series Anomaly Detection	Monitors sensor and infrastructure data	Enables predictive maintenance for branches, ATMs, and data centers

6. System Failures	Graph Convolutional Networks	Maps IT dependencies	Models software service dependencies to identify single points of failure and outage contagion risks
7. Process Management	AI-Driven Process Mining	Analyzes core banking event logs	Identifies workflow bypasses, unauthorized overrides, and ledger posting errors in real time

4. The Proposed Five-Layer Composite AI Architecture

To enable continuous, real-time risk mitigation, we propose a standardized, five-layer system architecture designed for scalable deployment within a bank's IT ecosystem (Gartner, 2023). Think of this architecture as a series of filters, each adding a layer of intelligence to the raw data flowing through the bank's systems.

Layer 1: Data Ingestion Layer

This layer is the "eyes and ears" of the system. It continuously captures data from multiple operational sources across the bank, including:

- Core banking transactions (OLTP systems)
- Security access logs
- Internal audit reports
- Customer complaint channels
- Historical RCSA checklists
- Employee records
- IT system logs

The key principle is **comprehensive data collection**—capturing everything that might contain signals of operational risk, rather than relying on limited samples (Basel Committee on Banking Supervision, 2021).

Layer 2: Preprocessing and Feature Engineering Layer

Raw data is messy and inconsistent. This layer cleans and organizes the data so the AI engines can work with it effectively. It:

- Extracts structured event logs for process mining
- Converts text documents into numerical representations (embeddings) using BERT
- Constructs relationship matrices representing the bank's organizational structure
- Standardizes data formats across different source systems
- Handles missing data and resolves inconsistencies

This layer is critical because AI models are only as good as the data they receive. Poor data quality leads to poor insights (Devlin et al., 2019; Hamilton et al., 2017).

Layer 3: Composite AI Engine Layer

This is the analytical core of the framework. It runs all four AI engines in parallel:

1. **Anomaly detection** (Isolation Forest and Deep Autoencoders)
2. **Text analysis** (BERT/NLP)
3. **Network analysis** (Graph Neural Networks)
4. **Process analysis** (Process Mining)

Each engine analyzes the data from its own perspective, and their outputs are combined into a multi-dimensional risk signature. This **Composite AI** approach is more powerful than any single technique because different types of risks are best detected by different methods (Gartner, 2023). Fraud might be detected by network analysis, process deviations by process mining, and emerging risks by anomaly detection.

Layer 4: Dynamic KRI and Decision Support Layer

Rather than relying on rigid, manually set thresholds, this layer dynamically updates risk indicators based on the continuous stream of insights from Layer 3. Real-time alerts are prioritized and routed to human risk analysts through a central dashboard.

What makes this "dynamic"? Traditional KRIs are static—they might be reviewed quarterly or annually. This system updates continuously, adapting to changing conditions. If the system detects that certain patterns are becoming more common, it can automatically adjust thresholds to maintain an appropriate balance between catching real risks and avoiding false alarms.

Layer 5: Explainable AI (XAI) Layer

This layer addresses one of the biggest concerns about AI in banking: the "black-box" problem. When an AI system makes a decision, how can we understand why?

Using techniques like SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations), this layer provides human-readable explanations for every alert (Lundberg & Lee, 2017; Ribeiro et al., 2016). It shows exactly which factors contributed to a risk score: an unusual login time, an atypical transaction amount, a skipped approval step, or an unusual connection between entities.

Why this matters: Regulators require banks to explain their decisions. If a bank cannot explain why its AI system flagged a transaction as suspicious, it may face regulatory challenges. XAI provides the transparency needed for regulatory compliance and internal audit (Sudjianto et al., 2020; Fritz-Morgenthal et al., 2022).

5. Practical Benefits of the Composite AI Framework

5.1. From Sampling to Full Coverage

Traditional risk management relies on sampling—reviewing a small percentage of transactions and hoping to catch problems. This is like checking a few cars on a highway and assuming all drivers are following the rules. The Composite AI framework monitors **100% of transactions and processes continuously**, catching issues that sampling would miss.

5.2. From Reactive to Proactive

Traditional tools only detect problems after they have occurred and losses have materialized. The AI framework detects **anomalies and deviations in real time**, often before any loss occurs. This shifts the bank from a reactive stance (responding to losses) to a proactive one (preventing losses).

5.3. From Rigid Rules to Adaptive Intelligence

Traditional systems rely on rules that must be manually defined and updated. The AI framework **learns from data** and adapts to changing conditions. It can detect new types of fraud or errors that have never been seen before, without requiring anyone to write a new rule.

5.4. From Ignoring Text to Reading Everything

Traditional risk management largely ignores unstructured text data. The AI framework **reads and analyzes** audit reports, customer complaints, regulatory documents, and employee communications, extracting valuable risk signals that would otherwise be missed.

5.5. From Individual Transactions to Connected Networks

Traditional analysis treats each transaction in isolation. The AI framework **sees the connections** between transactions, people, and systems, identifying collusion networks and systemic risks that would be invisible to traditional analysis.

6. Implementation Challenges and AI Governance

Implementing a Composite AI framework within a heavily regulated banking sector introduces critical organizational, technological, and regulatory challenges that require structured AI governance (Financial Stability Board, 2020; Bank for International Settlements, 2024).

6.1. Regulatory Acceptance and Explainable AI

Financial regulators, including central banks and national supervisory authorities, enforce strict standards on algorithmic models used for capital allocation and risk monitoring (Bank for International Settlements, 2024). The use of complex "black-box" deep learning models can lead to regulatory rejection if the internal decision-making process cannot be explained.

To gain regulatory approval, banks must integrate explainability directly into their risk management workflows (Sudjianto et al., 2020). By utilizing SHAP and LIME, the framework provides human risk analysts and external regulators with a clear, auditable explanation of every risk score, detailing exactly how each feature influenced the model's decision (Lundberg & Lee, 2017; Ribeiro et al., 2016).

6.2. Data Governance and Quality Assurance

The accuracy of the Composite AI framework is highly dependent on the quality and integrity of its input data. A key technical prerequisite is the standardization of core banking event logs (van der Aalst, 2016). If event logs contain missing timestamps, duplicate entries, or inconsistent user IDs, process mining algorithms will produce highly distorted process maps, leading to excessive false positive alerts (Dumas et al., 2018).

Banking institutions must establish robust data governance policies that ensure data cleanliness, completeness, and real-time synchronization across all core systems. This includes:

- **Data standards.** Consistent formats for dates, identifiers, and other critical fields.
- **Data quality monitoring.** Automated checks for missing or inconsistent data.
- **Data lineage.** Understanding where data comes from and how it has been transformed.
- **Master data management.** Ensuring consistent definitions of key entities (customers, accounts, products).

6.3. Data Privacy and Federated Learning

Collecting and analyzing employee behavior and customer transaction patterns raises significant data privacy concerns, particularly under strict regulations such as the European Union's General Data Protection Regulation (GDPR) (EBA, 2020). To mitigate these privacy risks, banks can deploy **Federated Learning** techniques (Goodfellow et al., 2016).

Under a Federated Learning architecture, machine learning models are trained locally at individual branch servers or distinct database nodes. The central risk management unit only aggregates the mathematical model parameters (gradients) rather than consolidating raw, sensitive customer or employee personal data. This ensures compliance with global data protection standards while still benefiting from collective learning across the organization.

6.4. Model Drift and Continuous Life-Cycle Management

Financial environments and fraudulent behaviors are highly dynamic. Over time, models trained on static historical data will experience **model drift**—a degradation in predictive accuracy due to changing operational environments and evolving fraud patterns (Fritz-Morgenthal et al., 2022).

To prevent model obsolescence, banks must adopt a continuous model lifecycle management framework. This requires:

- **Continuous performance monitoring.** Tracking model accuracy and alerting when performance degrades.
- **Automated testing.** Regularly validating models against new data.
- **Scheduled retraining.** Updating models with fresh data to keep them current.
- **Version control.** Tracking model versions to enable rollback if issues are discovered.

6.5. AI Literacy and Organizational Culture

Based on current industry trends, most banks are integrating AI risks within their existing operational risk taxonomies rather than creating a new standalone risk type (Deloitte & ORX, 2025). However, unique risks arising from AI such as hallucination and prompt injection do not naturally fall within traditional risk categories (Deloitte & ORX, 2025).

Governance and accountability need to be strengthened to ensure banks create and foster an AI risk-aware culture through clear risk leadership. This includes investing in key "no regret" actions such as **AI literacy at all levels** of the organization (Deloitte & ORX, 2025). Everyone from the board of directors to front-line staff needs to understand both the opportunities and risks of AI.

7. Conclusions and Future Research Directions

Traditional, retrospective operational risk management systems are increasingly inadequate in the face of modern, high-speed digital banking. The three pillars of traditional ORM—loss event databases, periodic self-assessments, and static key risk indicators—are simply too slow, too limited in scope, and too reliant on human judgment to effectively manage the complex, dynamic risks of contemporary banking.

This paper has proposed a comprehensive, multi-dimensional Composite AI framework that integrates four complementary AI methodologies:

1. **Unsupervised anomaly detection** (Isolation Forest and Deep Autoencoders) to find unknown risks
2. **Natural Language Processing** (BERT) to extract insights from unstructured text
3. **Graph Neural Networks** (GCN) to detect collusion and systemic connections
4. **Process mining** to identify deviations from intended workflows

These four techniques work together to provide real-time, continuous, and explainable operational risk monitoring (Gartner, 2023). The framework also incorporates an Explainable AI layer using SHAP and LIME to address the "black-box" concern that has historically limited AI adoption in regulated financial institutions.

For banking executives, transitioning to an AI-driven ORM architecture is no longer an optional luxury but a critical strategic requirement to ensure operational resilience and competitive survival (Basel Committee on Banking Supervision, 2023; Deloitte & ORX, 2025). The evidence is clear: AI is transforming operational risk landscapes, amplifying some risks while providing powerful tools to mitigate others. Banks that fail to adopt AI-driven risk management will not only miss opportunities for improved efficiency and risk reduction but may also face increased operational losses as their competitors leverage more sophisticated capabilities.

For regulators, there is a need to update supervisory frameworks to accommodate advanced AI risk monitoring, providing clear guidelines for the validation of explainable AI models (Bank for International Settlements, 2024). The development of AI-ready operational risk management frameworks will not be straightforward, although banks can leverage their existing strong foundations that have been developed over many years.

For future research, we identify two promising areas of development:

Reinforcement Learning for Automated Response. The integration of Reinforcement Learning models could enable autonomous decision-making, allowing the system to automatically block highly suspicious transactions or freeze compromised user credentials in real time (Goodfellow et al., 2016). This would represent a significant step beyond detection toward automated prevention.

Large Language Models for Regulatory Technology. Exploring the deployment of Large Language Models (LLMs) to automatically convert new regulatory circulars into executable compliance rules would represent a significant leap forward in regulatory technology (RegTech) (Financial Stability Board, 2020). This could dramatically reduce the cost and latency of regulatory compliance.

Finally, researchers should continue to investigate the relationship between AI investment and operational losses identified by McLemore and Mihov (2024). Understanding why AI investments sometimes increase

operational losses—and how to prevent this—will be critical for ensuring that AI adoption delivers its promised benefits without creating new vulnerabilities.

8. References

Arner, D. W., Barberis, J., & Buckley, R. P. (2017). The evolution of Fintech: A new post-crisis paradigm? *Georgetown Journal of International Law*, 47(4), 1271-1319.

Bank for International Settlements (BIS). (2024). Artificial intelligence in the financial sector: recent developments and main challenges. BIS Basel Reports.

Basel Committee on Banking Supervision (BCBS). (2021). *Principles for the Sound Management of Operational Risk*. Bank for International Settlements (BIS).

Basel Committee on Banking Supervision (BCBS). (2023). *Principles for Operational Resilience*. Bank for International Settlements (BIS).

de Medeiros, A. K. A., van der Aalst, W. M. P., & Weijters, A. J. M. M. (2008). Quantifying process conformance: A framework. *Cooperative Information Systems*, 17(3), 365-404.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT*, 1(1), 4171-4186.

Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of Process Mining*. Springer.

European Banking Authority (EBA). (2020). *Report on Big Data and Advanced Analytics*. EBA Reports.

Financial Stability Board (FSB). (2020). *Artificial Intelligence and Machine Learning in Financial Services: Market developments and financial stability implications*. FSB Publications.

Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, and responsible AI. *Frontiers in Artificial Intelligence*, 5, 849129.

Gartner. (2023). *Top Strategic Technology Trends for 2024: Composite AI*. Gartner Research.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations (ICLR)*.

Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Journal of Financial Regulation and Compliance*, 27(1), 99-136.

Liu, F. T., Ting, K. M., & Zhou, Z. H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1), 1-39.

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.

McLemore, P., & Mihov, A. (2024). AI and Operational Losses: Evidence from U.S. Bank Holding Companies. *Review of Corporate Finance Studies*, Accepted.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

Sudjianto, A., Knauth, W., Singh, R., Yang, Z., & Zhang, A. (2020). Explainable AI in finance. *arXiv preprint arXiv:2011.04271*.

van der Aalst, W. M. P. (2016). *Process Mining: Data Science in Action*. Springer, Berlin, Heidelberg.