

# Artificial Intelligence and Operational Risk Management: Current Applications, Challenges, and Research Gaps

Brijesh Shekhawat

Project Co Ordinator, AI and ML, Osmania Institute of Technology and Management, Hyderabad, Telangana

---

## Abstract

Operational risk — exposure to loss from failed internal processes, systems, people, or external events — has become one of the most active application areas for artificial intelligence (AI) in enterprise and financial risk management. This review examines current AI applications, persistent implementation challenges, and open research gaps in operational risk management. Industry survey evidence indicates that four in five banks now deploy AI for operational-risk management, with cyber risk the fastest-growing use case, even as governance concerns and unresolved questions about return on investment temper further deployment (Risk.net, 2026). Academic and applied literature documents AI's core contributions through anomaly detection and predictive analytics (Leo et al., 2019; Heß & Damásio, 2025), AI for IT operations (AIOps) for automated root-cause analysis and log-anomaly detection (Notaro et al., 2021), and enterprise-architecture approaches that embed operational-risk controls directly into configurable workflows (Basireddy, 2022b) and metadata-driven platforms (Basireddy, 2022a). The review further considers real-time data-processing requirements for operational-risk detection (Sannidhanam, 2021; Event Streaming Architectures for High-Volume Transaction Processing, 2022) and the emerging role of autonomous AI agents in continuous infrastructure risk monitoring (Sannidhanam, 2025). Despite substantial adoption, the review identifies persistent challenges — data quality, model interpretability, and a widening gap between AI deployment speed and organizational governance capacity (Risk.net, 2026) — and research gaps including limited empirical evaluation of AIOps effectiveness in production environments, underdeveloped frameworks for governing autonomous risk-remediation agents, and insufficient integration between operational-risk AI models and broader enterprise risk management architecture.

**Keywords:** Operational risk; artificial intelligence; AIOps; anomaly detection; enterprise architecture; AI governance; cyber risk; bankin

---

## Introduction

Operational risk — defined by regulators and risk-management practitioners as the risk of loss resulting from inadequate or failed internal processes, people, and systems, or from external events — has traditionally been the least quantifiable of the major risk categories facing financial institutions and large enterprises, in contrast to the more statistically tractable domains of credit and market risk. This has made operational risk a particularly attractive target for artificial intelligence (AI) applications, which are well suited to detecting anomalous patterns across the large, heterogeneous, and often unstructured datasets that operational-risk events generate — transaction logs, system telemetry, employee activity records, and vendor and third-party data. Recent industry survey

evidence indicates that AI adoption for operational-risk management has become mainstream practice rather than an emerging trend: Risk.net's 2026 Op Risk Benchmarking study, drawing on record participation from 61 banks globally, found that four in five institutions now deploy AI to manage operational risks, with cyber risk the fastest-growing use case within this category (Risk.net, 2026).

This review examines current applications of AI in operational risk management, the implementation and governance challenges these applications raise, and the research gaps that remain in the academic and applied literature. It draws on systematic reviews of machine learning in banking risk management (Heß & Damásio, 2025; Leo et al., 2019), applied research on AI for IT operations, or AIOps (Notaro

et al., 2021), industry survey data (Risk.net, 2026; Cambridge Centre for Alternative Finance, 2026), and applied enterprise-architecture research addressing how operational-risk controls are embedded into enterprise systems (Basireddy, 2022a, 2022b) and real-time data infrastructure (Sannidhanam, 2021, 2025; Event Streaming Architectures for High-Volume Transaction Processing, 2022).

## Current Applications of AI in Operational Risk Management

### Anomaly Detection in Transactions and Processes

The most widely deployed AI application in operational risk management is anomaly detection: identifying transactions, process executions, or user behaviors that deviate from established norms and may indicate error, fraud, or control failure. Systematic reviews of machine learning in banking risk management find this application area adequately covered in the academic literature, with the predictive and pattern-recognition capabilities of neural-network and ensemble methods well suited to the high-volume, largely structured data that operational and transaction-monitoring systems generate (Heß & Damásio, 2025; Leo et al., 2019). Applied industry evidence reinforces this, reporting substantial reductions in false-positive rates and improved detection accuracy when AI-driven monitoring is applied to anti-money-laundering and transaction-monitoring workflows.

### AI for IT Operations (AIOps)

A closely related application area is AIOps — the use of AI and machine learning to enhance IT operations management, including automated anomaly detection in system logs and automated root-cause analysis of service incidents. Applied research demonstrates hybrid frameworks combining unsupervised techniques such as principal component analysis with artificial neural networks to substantially reduce pseudo-positive anomaly detections in log-analysis pipelines, tested against both simulated and real-world datasets (Notaro et al., 2021). AIOps is particularly relevant to operational risk because IT-system failures and service disruptions constitute a major sub-category of operational-risk events, and

AIOps techniques directly target the detection and resolution speed of these failures, extending naturally into the broader operational-risk-management function.

### Cyber and Fraud Risk

Cyber risk has emerged as the fastest-growing AI application within operational risk management, according to recent industry survey evidence, with banks increasingly applying algorithmic tools to threat detection, anomaly monitoring, and fraud alerting (Risk.net, 2026). This trend is corroborated by broader financial-services survey research, which finds rapid AI deployment across the sector generally, alongside a recurring theme that AI deployment speed is outpacing institutional governance capacity (Cambridge Centre for Alternative Finance, 2026, as cited in Risk.net, 2026).

### Workflow-Embedded Operational Controls

Beyond detection-focused applications, applied enterprise-architecture research addresses how AI-informed operational-risk controls can be embedded directly into business-process workflows rather than delivered as separate, after-the-fact monitoring reports. Basireddy (2022b) proposes a configurable enterprise workflow architecture that digitizes risk management by embedding risk-control logic within operational workflows themselves, allowing risk events to be identified and addressed as part of normal process execution. This approach directly targets a limitation common to detection-only AI applications: identifying an operational-risk event is of limited value if the organization's workflow architecture cannot translate that detection into a timely operational response.

### Configurable, Metadata-Driven Platforms

As operational-risk requirements evolve — driven by changing regulations, business processes, and emerging risk types — static, purpose-built AI systems face growing maintenance and adaptation costs. Basireddy (2022a) proposes a metadata-centric platform architecture supporting dynamic configuration and intelligent process automation, allowing organizations to adjust operational-risk controls and automation logic without extensive custom redevelopment, an increasingly important

capability given the pace at which new operational-risk types (such as those associated with generative AI itself) are emerging.

Real-Time Infrastructure and Autonomous Monitoring

Operational-risk events, particularly those involving IT infrastructure or high-volume transaction processing, often require detection and response times that batch-oriented monitoring cannot support. Applied research on event streaming architectures for high-volume transaction processing (2022) and real-time claims processing using event-driven architectures (Sannidhanam, 2021) demonstrates how continuous, event-driven data processing supports lower-latency operational-risk detection. Most recently, autonomous AI agents for cloud infrastructure operations extend this further by enabling continuous monitoring paired with autonomous remediation action, representing the leading edge of AI application in operational-risk management, where the AI system not only detects but directly responds to operational anomalies (Sannidhanam, 2025). Complementary applied research on AI-enhanced IoT data analytics illustrates a parallel real-time monitoring application, integrating AI pattern recognition with data from Internet of Things devices to assess operational and transaction-related risk in banking operations (Thirumagal et al., 2024).

Evidence Snapshot: AI Applications in Operational Risk Management

Table 1 summarizes the principal AI application areas in operational risk management reviewed in this article.

Implementation and Governance Challenges

The Governance-Deployment Gap

Perhaps the most significant challenge identified in current evidence is a structural gap between the pace of AI deployment and the pace of governance development. Industry survey research finds that while AI deployment for operational-risk management has become widespread, governance

Table 1: AI applications in operational risk management.

| Application Area                          | AI Technique                                                                                      | Evidence / Source                                                                                                                                                                                 |
|-------------------------------------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Transaction and process anomaly detection | Unsupervised and supervised machine learning on operational telemetry and transaction logs        | Adequately covered in systematic banking-risk reviews (Heß & Damásio, 2025; Leo et al., 2019)                                                                                                     |
| IT operations / AIOps                     | Log anomaly detection combining PCA and artificial neural networks; automated root-cause analysis | Hybrid unsupervised frameworks reduce false-positive anomaly detections in production log data (Notaro et al., 2021; related AIOps literature)                                                    |
| Cyber and fraud risk                      | Algorithmic threat detection, anomaly monitoring, fraud alerting                                  | Fastest-growing AI use case within operational risk; 80% of surveyed banks deploy AI for operational risk overall (Risk.net, 2026)                                                                |
| Workflow-embedded controls                | Configurable enterprise workflow architecture                                                     | Risk-control logic embedded directly into operational processes rather than after-the-fact reporting (Basireddy, 2022b)                                                                           |
| Platform configurability                  | Metadata-centric, dynamically configurable platform architecture                                  | Enables reconfiguration of operational-risk controls without custom redevelopment (Basireddy, 2022a)                                                                                              |
| Real-time transaction processing          | Event-driven and event-streaming architectures                                                    | Supports low-latency detection of operational-risk events in high-volume transaction environments (Sannidhanam, 2021; Event Streaming Architectures for High-Volume Transaction Processing, 2022) |
| Autonomous infrastructure monitoring      | Autonomous AI agents for continuous monitoring and remediation                                    | Extends operational-risk management into self-directed infrastructure operations (Sannidhanam, 2025)                                                                                              |
| IoT-enabled operational monitoring        | AI-enhanced IoT data analytics                                                                    | Applied to real-time risk assessment of banking operations using IoT device data (Thirumagal et al., 2024)                                                                                        |

concerns and unresolved questions about return on investment continue to temper deeper deployment, and that AI deployment is generally outpacing institutional governance capacity across financial

services (Risk.net, 2026). This is consistent with academic caution that formal risk-management framework adoption does not, by itself, guarantee improved governance effectiveness, and that implementation quality remains decisive.

### **Data Quality and Interpretability**

Applied AIOps research identifies data quality — noisy, inconsistent, or incomplete log and telemetry data — as a persistent practical obstacle to reliable AI-based operational-risk detection, alongside challenges of user adoption and trust in automated anomaly-detection outputs. As with other AI risk-management applications, the most accurate detection models are often the least interpretable, creating friction with the auditability and explainability expectations of both internal risk functions and external regulators (Mayenberger, 2019).

### **Regulatory Expectations**

Operational-resilience regulatory expectations, including those articulated by the Basel Committee's Principles of Operational Resilience, increasingly require institutions to demonstrate not only that operational risks are being detected but that detection and response processes are documented, auditable, and subject to appropriate governance — a requirement that AI-based detection systems must be explicitly designed to satisfy rather than treated as a secondary consideration.

### **Autonomous Action and Control Boundaries**

As AI systems move from detection toward autonomous remediation — exemplified by autonomous AI agents for infrastructure operations (Sannidhanam, 2025) — organizations face a governance challenge that existing operational-risk frameworks, designed primarily around human-in-the-loop decision-making, do not fully address: defining the appropriate scope, oversight mechanisms, and escalation procedures for AI systems capable of taking direct remediation action within enterprise infrastructure.

### **Research Gaps**

Empirical evaluation of AIOps effectiveness in production environments remains comparatively limited, with much of the applied literature

emphasizing algorithmic performance in controlled or simulated settings rather than measured operational-risk-reduction outcomes in live deployments (Notaro et al., 2021).

Liquidity- and operational-risk interaction effects are underexplored: operational failures can trigger or compound liquidity and market-risk events, yet most AI models continue to be developed and evaluated within a single risk-category silo (Heß & Damásio, 2025).

Governance frameworks specific to autonomous, action-taking AI agents in operational-risk contexts remain underdeveloped relative to the pace of deployment of such systems (Sannidhanam, 2025).

Limited academic research directly measures the operational-risk impact of configurable, metadata-driven enterprise-architecture approaches (Basireddy, 2022a, 2022b), which remain more prominent in applied and industry literature than in peer-reviewed academic evaluation.

The relationship between AI deployment intensity and actual operational-risk-loss outcomes — as opposed to detection-accuracy metrics alone — remains insufficiently studied at the institutional level, despite high reported AI-adoption rates (Risk.net, 2026).

Cross-sector comparative research remains limited: most operational-risk AI literature concentrates on banking, leaving less-studied sectors — including healthcare, telecommunications, and industrial operations, where AIOps techniques are also applied — comparatively underexamined for operational-risk-specific outcomes.

### **Conclusion**

Artificial intelligence has moved from an emerging capability to mainstream practice in operational risk management, with the substantial majority of surveyed banks now deploying AI-based tools for this purpose and cyber risk emerging as the fastest-growing application area. The evidence reviewed here shows meaningful, well-documented applications spanning anomaly detection, AIOps, workflow-embedded controls, configurable enterprise platforms, and increasingly autonomous infrastructure monitoring. At the same time, the literature and industry survey

evidence converge on a consistent concern: AI deployment in operational risk management is currently outpacing the governance, interpretability, and empirical-evaluation infrastructure needed to manage it responsibly. Closing this gap — through more rigorous empirical evaluation of AI-based operational-risk tools in production settings, governance frameworks specifically designed for autonomous AI action, and stronger integration between AI models and enterprise-architecture platforms — represents the central research and practice agenda for the continued maturation of AI in operational risk management.

## References

1. Basireddy, S. (2022a). Pioneering a future-ready metadata-centric platform architecture for dynamic configuration intelligent process automation operational agility and sustainable enterprise modernization. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 14(03), 415-429. <https://doi.org/10.18090/samriddhi.v14i03.27>
2. Basireddy, S. R. (2022b). Digitizing risk management through configurable enterprise workflow architecture. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 14(4), 231-239.
3. Event Streaming Architectures for High-Volume Transaction Processing. (2022). International Journal of Advance Industrial Engineering, 10(01), 1-8.
4. Heß, V. L., & Damásio, B. (2025). Machine learning in banking risk management: Mapping a decade of evolution. International Journal of Information Management Data Insights, 5(1), Article 100324. <https://doi.org/10.1016/j.jjime.2025.100324>
5. Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. Risks, 7(1), 29. <https://doi.org/10.3390/risks7010029>
6. Mayenberger, D. (2019). How to overcome modelling and model risk management challenges with artificial intelligence and machine learning. Journal of Risk Management in Financial Institutions, 12(3), 241-255.
7. Notaro, P., Cardoso, J., & Gerndt, M. (2021). A survey of AIOps methods for failure management. ACM Transactions on Intelligent Systems and Technology, 12(6), 1-45. <https://doi.org/10.1145/3483424>
8. Risk.net. (2026). 2026 Op Risk Benchmarking: Banks adopt AI to manage operational risks. Infopro Digital Risk (IPDR).
9. Sannidhanam, A. H. (2021). Real-time claims processing using event-driven architectures. International Journal of Technology, Management and Humanities, 7(01), 36-50. <https://doi.org/10.21590/07.01.02>
10. Sannidhanam, A. H. (2025). Autonomous AI agents for cloud infrastructure operations. Journal of Contemporary Science and Technology Management, 1(01), 83-100.
11. Thirumagal, P. G., Vaddepalli, S., Das, T., Das, S., Madem, S., & Immaculate, P. S. (2024). AI-enhanced IoT data analytics for risk management in banking operations. In Proceedings of the 5th International Conference on Recent Trends in Computer Science and Technology (ICRTCST 2024) (pp. 177-181). IEEE.