Economic Sciences
https://economic-sciences.com
ES (2026) 22, 5(S), 245 -250 | ISSN:1505-4683

The Evolution of AI-Driven Predictive Modelling in Fintech: A Systematic Survey of Risk, Fraud, and Cybersecurity

Deepshikha Aggarwal
Department of Information Technology, Jagan Institute of Management Studies, Delhi, India
deepshikha.aggarwal@jimsindia.org
ORCID: https://orcid.org/0000-0002-2782-0733

Abstract

As the financial services sector transitions to real-time, cross-border digital architectures, traditional risk management frameworks face unprecedented challenges from increasing transaction velocity and sophisticated, AI-driven threats. This paper provides a comprehensive systematic survey of the academic literature published between 2021 and 2026, focusing on the convergence of Artificial Intelligence (AI) in risk management, fraud detection, and cybersecurity.Through a PRISMA-guided review of over 30 core studies, we categorize the evolution of methodologies from static statistical models to dynamic, multimodal intelligence systems. Key findings indicate a significant shift toward the use of Graph Neural Networks (GNNs) for detecting systemic contagion and the integration of Large Language Models (LLMs) for proactive market sentiment analysis. Furthermore, this survey explores the critical "Accuracy-Explainability" trade-off, analysing how Explainable AI (XAI) and Federated Learning have emerged as essential tools for regulatory compliance under the EU AI Act and DORA frameworks.The analysis reveals that while AI significantly enhances predictive accuracy, it also introduces new vulnerabilities, specifically regarding adversarial machine learning and synthetic identity fraud. We conclude by identifying persistent research gaps, including the need for quantum-resistant security protocols and more robust defences against generative adversarial deception. This study serves as a strategic roadmap for researchers and practitioners navigating the "arms race" between defensive AI and industrialized financial crime.

Keywords

Primary: Artificial Intelligence, Fintech, Predictive Risk Modelling, Fraud Detection, Cybersecurity.

1. Introduction

The global financial landscape is currently undergoing a "Real-Time Revolution." As transactions move from batch processing to instantaneous, cross-border digital exchanges, the window for detecting risk has shrunk from days to milliseconds. Traditional risk management frameworks, which relied heavily on historical statistical inference and static rule-based systems, are increasingly inadequate against the velocity of modern financial threats (Zhang et al., 2024).

The convergence of Artificial Intelligence (AI) and Financial Technology (Fintech) has shifted the paradigm from reactive mitigation to proactive intelligence. While early AI applications in finance were limited to narrow tasks like credit scoring, the current academic frontier focuses on **Predictive** Risk Modelling—the ability to anticipate market volatility, systemic contagion, and sophisticated cyber-fraud before they manifest.

The historical trajectory of risk assessment began with the Altman Z-score (Altman, 1968) and early credit scoring models (West, 2000), which prioritized linear interpretability. However, the 2008 financial crisis exposed the fragility of these models in the face of systemic "contagion" (Haldane & May, 2011). The subsequent decade saw a rapid shift toward ensemble learning, with Random Forests and XGBoost becoming the industry standard for tabular fraud detection (Jurgovsky et al., 2018; Chen et al., 2020). Today, the focus has shifted again—this time toward adversarial resilience as hackers use the same AI tools to probe for model vulnerabilities (Goodfellow et al., 2015; Thompson, 2026).

This survey paper systematically examines the recent academic literature (2021–2026) regarding AI's role in three critical pillars:

Economic Sciences
https://economic-sciences.com
ES (2026) 22, 5(S), 245 -250 | ISSN:1505-4683

1. **Risk Management:** Predictive modelling for credit and market stability.
2. **Fraud Detection:** Identifying synthetic identities and anomalous transaction patterns.
3. **Cybersecurity:** Defending the financial infrastructure against AI-powered adversarial attacks.

The objective of this study is to categorize the dominant AI architectures, evaluate their performance metrics, and identify the persistent "open challenges" that define the current research gap.

2. Literature Review

2.1 From Tabular to Multimodal Risk Intelligence

Historically, academic research focused on **Structured Data Models**. Standard machine learning algorithms, such as Extreme Gradient Boosting (XGBoost) and Random Forests, dominated the literature due to their high performance on tabular financial records (Chen & Li, 2022). However, recent studies argue that these models suffer from "contextual blindness"—they cannot account for the narrative-driven nature of modern markets.

The "Third Wave" of research (2024–2026) emphasizes Multimodal Learning. Researchers like Nguyen et al. (2025) have demonstrated that combining numerical transaction data with unstructured data—such as social media sentiment and geopolitical news via Large Language Models (LLMs)—improves market risk prediction accuracy by over 15% compared to uni-modal models.

2.2 Graph Neural Networks (GNNs) and Contagion Risk

A significant trend in the literature is the application of Graph Neural Networks (GNNs) to address systemic risk and money laundering. Unlike traditional models that treat transactions as isolated events, GNNs treat the financial system as a complex web.

Recent work by Al-Mansour (2025) highlights how GNNs can detect "smurfing" (breaking large sums of money into small transactions) by analysing the topological structure of transaction networks. This research suggests that the *relationship* between entities is often more predictive of risk than the *attributes* of the entities themselves.

2.3 The Rise of Explainable AI (XAI) in RegTech

As AI models become more complex (e.g., Deep Neural Networks), they often become "black boxes." This has created a friction point with financial regulators who require "Right to Explanation" under frameworks like the AI Act of 2024 (European Parliament, 2024).

Academic literature has responded with a surge in Explainable AI (XAI) research. Studies by Miller and Gupta (2025) compare the efficacy of SHAP (SHapley Additive exPlanations) and LIME in credit risk scoring. Their findings suggest that while XAI adds a computational overhead, it is essential for identifying "Model Bias"—ensuring that predictive models do not inadvertently discriminate based on proxy variables for race or gender.

2.4 Adversarial Cybersecurity in Fintech

Finally, a critical sub-sector of the literature addresses the "AI Arms Race." As financial institutions deploy AI for defence, threat actors use Generative Adversarial Networks (GANs) to probe these systems for weaknesses. Research by Thompson (2026) explores "Adversarial Robustness," proposing models that are intentionally trained on "poisoned" data to better recognize the subtle deviations used by modern cyber-criminals to bypass biometric and behavioural security filters.

2.5 Privacy-Preserving Collaborative Learning A significant barrier in academic research has been the "Data Silo" problem. Financial institutions are legally barred from sharing raw data, which limits the training sets for AI models. Recent breakthroughs in Federated Learning (FL) (McMahan et al., 2017) and Differential Privacy (Abadi et al., 2016) have enabled "collaborative defence." Studies by Shamsabadi et al. (2025) and Alhchaimi (2024) demonstrate that FL-trained models can achieve accuracy parity with centralized models while maintaining strict GDPR compliance. Furthermore, the use of Generative Adversarial Networks (GANs) to create high-fidelity synthetic

Economic Sciences
https://economic-sciences.com
ES (2026) 22, 5(S), 245 -250 | ISSN:1505-4683

data has emerged as a key method for testing "stress scenarios" without compromising real user privacy (Tiwald et al., 2021; Dama, 2024).

3. Methodology: Systematic Review Process

To ensure a comprehensive and unbiased selection of literature, this survey follows the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines.

3.1 Search Strategy and Data Sources

The primary research was conducted across four major digital libraries: **IEEE Xplore, ACM Digital Library, ScienceDirect (Elsevier), and SSRN**. The search strings utilized Boolean operators to intersect the three core domains:

- ("Artificial Intelligence" OR "Machine Learning") AND ("Risk Management" OR "Fintech")
- ("Fraud Detection" OR "Anomaly Detection") AND "Deep Learning"
- ("Cybersecurity" OR "Adversarial AI") AND "Financial Services"

3.2 Inclusion and Exclusion Criteria

Studies were filtered based on the following criteria to maintain the survey's contemporary relevance:

1. **Temporal Scope:** Limited to peer-reviewed articles and conference proceedings published between January 2021 and April 2026.
2. **Technological Focus:** Papers must focus on *predictive* capabilities rather than purely descriptive or historical analysis.
3. **Application Domain:** Must be explicitly situated within the Fintech, Banking, or Decentralized Finance (DeFi) sectors.
4. **Exclusion:** Papers lacking empirical validation or those focusing solely on cryptocurrency price speculation (without a risk framework) were excluded.

4. Comparative Analysis of AI Architectures

This section synthesizes the findings from the surveyed literature, comparing the efficacy of various AI architectures across the pillars of risk, fraud, and cybersecurity. While accuracy remains a primary KPI, researchers have introduced nuanced metrics such as the Brier Score for calibration and Precision-Recall AUC for imbalanced datasets common in fraud (Kajal & Kaur, 2021; Lamgade, 2024). Studies in emerging markets show that trust in these metrics is a prerequisite for user adoption of mobile financial services (MDPI, 2025; Vishva, 2024).

4.1 Performance Comparison Table

Based on the aggregated results of the surveyed studies, the following table summarizes the strengths and weaknesses of the dominant architectures identified in the 2024–2026 research cycle.

Architecture	Primary Use Case	SOTA Performance Metric	Major Limitation
Graph Neural Networks (GNN)	Anti-Money Laundering (AML)	High "Community Detection" accuracy (89%+)	High computational latency for real-time edge use.
Transformers (LLM-based)	Market Risk & Sentiment	Superior F1-scores in unstructured data fusion.	Susceptibility to "Hallucination" in risk forecasting.
Federated Learning (FL)	Cross-Bank Fraud Detection	Privacy-preserving accuracy parity with centralized models.	Communication overhead and "Client Drift."
Ensemble (XGBoost/LightGBM)	Credit Scoring	Highest interpretability and speed for tabular data.	Poor handling of temporal/dynamic behaviour shifts.

4.2 The "Accuracy-Explainability" Trade-off

A recurring theme in recent academic papers (e.g., Wang & Zhao, 2025) is the Pareto Frontier of AI in finance. As models move toward Deep Learning to capture complex risk patterns, their "Transparency Score" decreases.

Economic Sciences
https://economic-sciences.com
ES (2026) 22, 5(S), 245 -250 | ISSN:1505-4683

Research indicates that while Deep Reinforcement Learning (DRL) is the most effective for dynamic portfolio risk management, its adoption in regulated banking is slower than Explainable Boosting Machines (EBM). This suggests that in the current academic climate, "Interpretability" is becoming as critical a performance metric as "Accuracy" (Miller & Gupta, 2025).

4.3 Resilience Against Adversarial Attacks

A critical finding in the 2025–2026 literature is the vulnerability of predictive models to Evasion Attacks. In several surveyed studies, researchers demonstrated that subtle "noise" injected into transaction metadata could cause a SOTA fraud detection model to misclassify a high-risk event as legitimate (Thompson, 2026). This has led to a new sub-field of study: Certified Robustness, where models are mathematically guaranteed to remain accurate within a specific "perturbation budget."

5. Open Challenges and Future Directions

Despite the significant advancements identified in Sections 3 and 4, the academic literature from 2024–2026 reveals critical "bottlenecks" that prevent the full autonomous operation of AI in Fintech.

5.1 The "Arms Race": Generative AI and Synthetic Deception

A primary challenge identified in recent studies is the industrialization of deception. Tresner (2026) observes that GenAI has moved beyond simple phishing to creating "synthetic identity infrastructure," including realistic identity documents and deepfake video for bypassing Know-Your-Customer (KYC) biometric gates.

- **Future Direction:** Research is shifting toward **"Liveness Detection 2.0,"** utilizing heart-rate monitoring via webcams and specialized AI that detects the subtle pixel-level "noise" inherent in deepfake generation (Seguin, 2026).

5.2 Agentic AI and Autonomous Risk Mitigation

One of the most significant trends in late 2025 research is the rise of Agentic AI—autonomous agents that don't just alert humans to risk but take preemptive action.

- **The Challenge:** Entrusting an AI agent to freeze global liquidity or block high-value corporate transfers introduces "Operational Fragility."
- **Future Direction:** Studies suggest a shift toward "Human-in-the-Loop-Enabled Autonomy," where AI handles 99% of low-level anomalies while using Predictive Uncertainty metrics to flag high-stakes decisions for human review (Citizens Bank Research, 2026).

5.3 Privacy-Preserving Collaborative Défense

Traditionally, banks have been hesitant to share fraud data due to competition and privacy laws like GDPR and DORA. However, isolated defences are increasingly insufficient against coordinated global attacks.

- **Future Direction:** Academic focus is heavy on Federated Learning (FL) and Confidential Computing. These technologies allow institutions to train a shared global "Threat Intelligence" model without ever sharing raw customer data (Gartner, 2026; Jack Henry, 2026).

6. Conclusion

The transition from traditional statistical auditing to AI-driven predictive modelling represents a seismic shift in the financial services sector. This survey has tracked the evolution of this field from the foundational linear models of the late 20th century (Altman, 1968) to the highly sophisticated, multimodal Transformer architectures and Graph Neural Networks that define the 2024–2026 research landscape (Nguyen et al., 2025; Al-Mansour, 2025).

6.1 Synthesis of Findings

Our analysis confirms that the "three-pillar" synergy of Risk Management, Fraud Detection, and Cybersecurity is no longer a theoretical goal but an operational necessity. Several critical themes emerged from the synthesized literature:

- **The Demise of Data Silos:** Through the advancement of Federated Learning (McMahan et al., 2017; Shamsabadi et al., 2025), the academic community has found a path to

reconcile the paradox between the need for big data and the strict mandates of privacy laws like GDPR and DORA.

- **The Explainability Mandate:** As models have moved toward "black-box" Deep Learning, the rise of XAI (Lundberg & Lee, 2017; Miller & Gupta, 2025) has bridged the gap between algorithmic performance and regulatory compliance.
- **The Proactive Intelligence Pivot:** The most significant trend in recent years is the move away from "lagging indicators" toward "leading indicators." By utilizing behavioural biometrics and high-frequency sentiment analysis, modern models can now forecast potential liquidity crises or fraud events before the final transaction is executed (Zhang et al., 2024).

6.2 The Persistent "Human-AI" Friction

Despite these breakthroughs, a core tension remains in the "Last Mile" of risk management. The literature suggests that while AI can automate 99% of anomaly detection, the final 1%—the high-stakes, "Black Swan" events—still requires human intuition (Haldane & May, 2011; Citizens Bank Research, 2026). The risk of Model Over-reliance is a growing concern in the academic sphere; if human analysts lose the ability to critically question AI outputs, the financial system may become more vulnerable to systemic, synchronized failures.

6.3 A Call for Future Research

Looking toward the 2030 horizon, this survey identifies three urgent areas for future academic inquiry:

1. **Adversarial Game Theory:** Developing models that can anticipate how human fraudsters will adapt their behaviour in response to new AI defences (Thompson, 2026; Tresner, 2026).
2. **Quantum-Resistant Risk Modelling:** As quantum computing approaches maturity, the cryptographic foundations of cybersecurity in fintech will require a complete overhaul, a topic currently under-represented in the literature.
3. **Cross-Sector Resilience:** Research must expand beyond the "Banking" silo to examine how AI-driven risk in the energy and supply chain sectors cascades into financial instability (Gartner, 2026).

In conclusion, AI is no longer a peripheral "efficiency tool" for financial institutions; it is the central nervous system of modern risk management. As we move into an era of "Industrialized Deception" powered by Generative AI, the resilience of the global economy will depend not on the volume of data we collect, but on the speed and transparency of the intelligence we derive from it.

References

1. **Abadi, M., et al. (2016).** Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC.*
2. **Akinsola, K. (2025).** Regulatory Compliance and Corporate Governance in Preventing Financial Misstatements. *Bostone Academic Press.*
3. **Alhchaimi, A. A. J. (2024).** Analyzing machine learning algorithms for cloud-based transaction fraud detection. *Wasit J. Comput. Math. Sci.*
4. **Al-Mansour, A. (2025).** Topological Financial Intelligence: GNNs in Anti-Money Laundering. *Journal of Financial Cyber-Security.*
5. **Altman, E. I. (1968).** Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy. *Journal of Finance.*
6. **Bamigboye, O. (2020).** Financial fraud prevention through strengthened corporate governance. *J. Law Glob. Policy.*
7. **Chen, H., & Li, X. (2022).** A Comparative Study of Ensemble Learning in Credit Scoring. *Fintech Research Quarterly.*
8. **Chen, T., et al. (2020).** XGBoost: A Scalable Tree Boosting System. *KDD Conference Proceedings.*
9. **Citizens Bank Research. (2026).** 2026 AI Trends in Financial Management: The Rise of Agentic AI. *Corporate Insights Division.*
10. **Dama, K. (2024).** Fraud Detection in Financial Transactions: A GAN-based Approach. *Kalasalingam Academy.*
11. **European Parliament. (2024).** The Artificial Intelligence Act: Implications for the Financial Sector.
12. **Gartner. (2026).** Predictive Analytics and Confidential Computing in Global Finance. *Research Report G-442.*

13. **Goodfellow, I. J., et al. (2015).** Explaining and Harnessing Adversarial Examples. *ICLR.*
14. **Gunning, D. (2017).** Explainable Artificial Intelligence (XAI). *DARPA Research.*
15. **Haldane, A. G., & May, R. M. (2011).** Systemic risk in banking ecosystems. *Nature.*
16. **Jack Henry & Associates. (2026).** The Structural Shift in Banking: Risk-Adaptive Authentication.
17. **Jurgovsky, J., et al. (2018).** Sequence classification for credit-card fraud detection. *Expert Syst. Appl.*
18. **Kajal, D., & Kaur, K. (2021).** Credit card fraud detection using imbalance resampling. *Int. J. Adv. Trends Comput. Sci.*
19. **Lamgade, P. (2024).** Addressing Class Imbalance in Financial Fraud Data. *Journal of Applied AI.*
20. **Lundberg, S., & Lee, S. (2017).** A Unified Approach to Interpreting Model Predictions. *NeurIPS.*
21. **McMahan, B., et al. (2017).** Communication-Efficient Learning of Deep Networks from Decentralized Data. *AISTATS.*
22. **MDPI Research. (2025).** AI-Driven Cybersecurity in Mobile Financial Services: Emerging Markets. *MDPI Global.*
23. **Miller, S., & Gupta, R. (2025).** Transparency vs. Accuracy: The XAI Trade-off in Regulatory Technology. *Int. J. AI & Law.*
24. **Nguyen, T., et al. (2025).** Multimodal Transformers for Proactive Market Risk Assessment. *IEEE Fintech.*
25. **Seguin, M. (2026).** Behavioral Signals in the Age of Manipulation. *Thomson Reuters Fintech Analysis.*
26. **Shamsabadi, A., et al. (2025).** Privacy-Preserving Machine Learning for Financial Risk Assessment. *IEEE Transactions.*
27. **Thompson, L. (2026).** Adversarial Resilience: Shielding AI-Driven Banks from Synthetic Fraud. *Cybersecurity Review.*
28. **Tiwald, P., et al. (2021).** Financial Risk Management and Explainable, Trustworthy AI. *PMC Research.*
29. **Tresner, K. (2026).** Industrialized Deception: The AI Threat Multiplier. *J. Fin. Crime & Tech.*
30. **Vishva, S. (2024).** Advanced Cybersecurity Measures and ML in Digitized Financial Systems. *Cyber Security Journal.*
31. **Wang, J., & Zhao, K. (2025).** The Pareto Frontier of AI in Fintech. *Journal of Computational Finance.*
32. **Zhang, Y., et al. (2024).** Beyond Hindsight: The Shift Toward Predictive Risk Intelligence. *Global Finance Journal.*
33. **Aggarwal, D., Sharma, D., & Saxena, A. B.** (2024). Smart education: an emerging teaching pedagogy for interactive and adaptive learning methods. Journal of Learning and Educational Policy, 44, 1-9.
34. **Aggarwal, D., Sharma, D., & Saxena, A. B.** (2024). Exploring the role of AI for enhancement of social media marketing. Journal of Media, Culture and Communication, 4(5), 1-11.
35. **Aggarwal, D.** (2019). Mobile technology adoption by Indian consumers. International Journal of Recent Technology and Engineering, 8(2), 892-899.
36. **D. Grover (2024)**, "The AI Assistant Revolution: Microsoft Copilot and The Future of Programming," Educational Administration Theory and Practice, vol. 30, no. 1.
37. **Lowe, D., & Galhotra, B. (2024).** A Model for Teaching and Evaluation of Programming Courses in Online mode. Journal of Informatics Education and Research, 4(1).