

Artificial Intelligence In Real-Time Payment Systems: A Bibliometric And Thematic Analysis Of Emerging Research Trends (2004–2026)

Vijaya Kumar*

Department of Commerce, Government First Grade college of Arts, Science and Commerce - SIRA, 572137

ABSTRACT

Real-time payment systems have transformed retail and wholesale finance by compressing payment initiation, authorization, clearing, and settlement into seconds or near-real time. The same speed that improves convenience and liquidity also narrows the window in which fraud, anomalies, cyber events, and financial-crime patterns can be detected. Artificial intelligence (AI) is therefore becoming a critical intelligence layer in payment infrastructures. This study maps the intellectual, thematic, and geographic development of research at the intersection of AI and real-time payments using a Scopus-derived corpus of 545 documents published across 402 sources from 2004 to 2026 and analyzed in Biblioshiny. Performance analysis, keyword co-occurrence, trend-topic analysis, thematic mapping, and collaboration analysis are combined to identify the field's knowledge structure and emerging research fronts. Publication activity accelerated sharply after 2023, reaching 179 documents in 2025, while conference papers account for 54.3% of the corpus, indicating a technically driven and rapidly evolving domain. The thematic map identifies two mature motor themes: learning systems–crime–fraud detection and artificial intelligence–blockchain–security. Graph neural networks and behavioral research occupy the lower-centrality, lower-density zone; their 2025–2026 temporal profile indicates that they are emerging rather than declining themes. India, the United States, and China dominate country-affiliation productivity, with India–United States collaboration forming the strongest international link. The study argues that the field is moving from transaction-level classification toward adaptive, network-aware, explainable, and privacy-preserving intelligence embedded in payment orchestration. A research agenda is proposed around temporal learning, graph analytics, federated models, explainability, behavioral scam detection, adversarial resilience, and AI-assisted liquidity management.

Keywords: artificial intelligence; real-time payments; instant payments; machine learning; fraud detection; graph neural networks; Biblioshiny; bibliometric analysis; thematic analysis.

INTRODUCTION

Fast or real-time payment systems (RTPS) have become a central component of digital financial infrastructure because they allow funds to be transferred and made available to recipients in real time or near real time, typically on a near-continuous basis. Their economic significance extends beyond speed. Fast-payment rails can lower transactional frictions, broaden participation in digital finance, and provide a foundation on which banks, fintech firms, merchants, governments, and households build new services (Frost et al., 2024; Cornelli et al., 2024). Yet instantaneity changes the risk architecture of payments. In slower systems, suspicious activity may be reviewed before final settlement or recovered

through operational delay. In real-time systems, authorization, fraud screening, compliance checks, and customer authentication must occur within a narrow decision window. A mistaken approval can become an irrevocable loss within seconds, while an overly conservative model can block legitimate transactions and undermine trust in the payment rail.

Artificial intelligence (AI) and machine learning (ML) offer a set of computational capabilities that are well matched to this environment: non-linear classification, anomaly detection, sequence learning, graph analytics, behavioral modeling, natural-language processing, reinforcement learning, and, increasingly, generative and agentic AI. Earlier payment-fraud research demonstrated that real-world

Relevant conflicts of interest/financial disclosures: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

usefulness depends not only on predictive accuracy but also on latency, class imbalance, evolving fraud vectors, data quality, and the business cost of false positives (Ryman-Tubb et al., 2018). Recent central-bank research extends the problem beyond card fraud. Layered ML architectures can identify unusual high-value payment activity in real time (Desai et al., 2024), network-wide analytics can reveal coordinated financial-crime patterns that are invisible to individual institutions (Bank for International Settlements [BIS] Innovation Hub, 2025), and generative AI agents are beginning to be tested for intraday liquidity and payment-prioritization decisions (Aldasoro & Desai, 2025). These developments indicate that AI is shifting from a peripheral fraud-detection tool to a broader operational and governance layer within payment infrastructures.

The scholarly literature, however, remains fragmented across computer science, cybersecurity, fintech, blockchain, banking, fraud analytics, Internet of Things (IoT), and payment-system research. Broad bibliometric studies of AI and ML in finance identify fraud, distress, forecasting, risk, and fintech as major clusters (Goodell et al., 2021; Pattnaik et al., 2024), while bibliometric work on explainable AI emphasizes interpretability, trust, risk assessment, and regulatory relevance (Chen et al., 2023). Those reviews are valuable but do not isolate the specific combination of AI and real-time payment environments, where low-latency decisions, rapid finality, network effects, interoperability, fraud migration, and operational resilience interact. Conversely, engineering studies often examine a single model, dataset, or payment problem without locating that contribution within the larger evolution of the field.

A bibliometric and thematic approach is useful because the field is expanding quickly and is distributed across disciplines. Bibliometrics can quantify scientific production and impact, while science mapping can expose conceptual relationships that are difficult to see through narrative reading alone (Donthu et al., 2021; Zupic & Čater, 2015). Co-word analysis is especially appropriate for an emerging interdisciplinary domain because repeated keyword combinations reveal how researchers connect technologies, risks, and applications over time

(Callon et al., 1991; Cobo et al., 2011). The present study therefore uses the Biblioshiny report generated from a Scopus-derived corpus to examine how research on AI in real-time payments has developed, which actors and sources structure the field, which topics have matured, and which research fronts are now emerging.

The study addresses four research questions. RQ1 asks how scientific production and citation impact have evolved in AI–real-time payment research. RQ2 asks which sources, authors, institutions, countries, and collaborations contribute most strongly to the field. RQ3 asks which conceptual themes organize the knowledge structure and how those themes have changed over time. RQ4 asks which emerging themes should shape the next research agenda for intelligent payment systems. By answering these questions, the article contributes a focused map of a rapidly forming research domain, distinguishes mature motor themes from new fronts, and develops a forward-looking framework that connects technical model development with payment-system requirements for speed, security, explainability, privacy, and resilience.

2. Conceptual Background: AI as an Intelligence Layer in Real-Time Payments

The defining characteristic of an RTPS is not merely faster messaging but the compression of the complete decision cycle around a transfer. Fast-payment systems generally aim to provide immediate or near-immediate availability of funds and broad operating hours, with design choices varying across jurisdictions (Committee on Payments and Market Infrastructures [CPMI], 2016; Frost et al., 2024). This architecture increases the value of automation because large transaction volumes must be evaluated continuously. It also heightens the cost of delayed intelligence: the interval between observing a suspicious event and taking preventive action may be shorter than the time required for traditional manual review.

Three analytically distinct layers of AI application can be identified. The first is transaction intelligence. Here, models score individual transactions using historical behavior, device or channel attributes, transaction sequences, and contextual features. Supervised classifiers remain useful when reliable

labels exist, but fraud detection often suffers from severe class imbalance, delayed labels, and changing attack patterns. Anomaly-detection and hybrid methods are therefore important for identifying previously unseen behaviors. Desai et al. (2024), for example, demonstrate a layered approach in which supervised learning first separates typical from unusual high-value payments and unsupervised learning then ranks anomalies among the unusual subset. The architecture illustrates a core RTPS design principle: computational resources and human attention should be concentrated on the smallest set of transactions that warrant intervention.

The second layer is network intelligence. Fraud and money laundering rarely operate as isolated transactions; they involve relationships among accounts, merchants, devices, mule networks, counterparties, and institutions. Graph neural networks (GNNs) can learn from these relational structures rather than treating each transaction as independent. A systematic review by Motie and Raahemi (2024) shows that GNN-based financial-fraud research is expanding because graph representations can capture patterns that conventional tabular models miss, while important gaps remain in unsupervised, dynamic, and multi-level anomaly detection. Project Hertha provides a payments-specific illustration: network-wide analytics applied to a large synthetic retail-payment ecosystem improved the detection of illicit accounts, with particularly strong gains for novel patterns (BIS Innovation Hub, 2025). Such evidence helps explain why GNN-related terms appear as an emerging frontier in the present bibliometric corpus.

The third layer is system intelligence. AI can support decisions about routing, liquidity, settlement priorities, operational anomalies, and capacity rather than only policing end-user transactions. Aldasoro and Desai (2025) experimentally evaluate a generative AI agent for intraday liquidity management in a wholesale payment setting and find that it can reproduce several prudential cash-management behaviors under simulated constraints. Reinforcement learning has also been applied to optimization problems involving blockchain and distributed infrastructures, including the highly cited work by Liu et al. (2019). In this layer, the AI system is no longer only a classifier; it becomes a decision-

support or decision-making component inside a complex financial market infrastructure. That transition magnifies governance requirements because model errors can affect liquidity, settlement queues, or systemic resilience.

Across all three layers, four constraints recur. First, latency matters: a model that is accurate but too slow may be unusable in a real-time authorization path. Second, fraud is non-stationary: adversaries adapt to controls, producing concept drift and novel behaviors. Third, models operate under asymmetric error costs. False negatives create losses and compliance exposure; false positives impose friction on legitimate users and can suppress adoption. Fourth, high-performing black-box models may conflict with requirements for explainability, auditability, contestability, and supervisory review. Research on explainable AI in finance accordingly identifies transparency and risk assessment as persistent themes (Chen et al., 2023). These constraints imply that the most valuable AI research for RTPS will not be defined by accuracy alone but by the joint optimization of detection quality, response time, interpretability, privacy, resilience, and operational cost.

3. Methodology

This study adopts bibliometric performance analysis and science mapping to examine the evolution and conceptual structure of research on AI in real-time payment systems. Bibliometric analysis is appropriate for large, heterogeneous literatures because it provides reproducible quantitative measures of productivity, influence, collaboration, and conceptual relatedness (Donthu et al., 2021; Zupic & Čater, 2015). The analysis was conducted from the supplied Biblioshiny workbook, which reports a Scopus-derived corpus and the associated Bibliometrix/Biblioshiny outputs. Bibliometrix is an R-based science-mapping environment designed for comprehensive bibliometric workflows, while Biblioshiny provides an interactive interface for comparable analyses and visualizations (Aria & Cuccurullo, 2017; Aria et al., 2026).

The verified corpus contains 545 documents published between 2004 and 2026 across 402 sources. The workbook records 55,452 references, 1,825 authors, 1,727 author keywords, and 3,136 Keywords

Plus terms. The average document age is 1.95 years and the mean citation count is 9.415 per document. International co-authorship is 25.32%, and the average number of co-authors per document is 3.78. Conference papers are the largest document type (296; 54.3%), followed by articles (158; 29.0%), book chapters (51; 9.4%), reviews (23; 4.2%), books (12; 2.2%), conference reviews (3; 0.6%), one data paper, and one letter. The composition is consistent with a technically oriented field in which conference dissemination often precedes journal consolidation.

The underlying search logic targeted the intersection of two concept families: AI-related terms (for example, artificial intelligence, machine learning, deep learning, neural networks, intelligent systems, AI-driven approaches, and predictive analytics) and fast-payment terms (for example, real-time payment, instant payment, faster payment, real-time transaction, and instant payment system). The supplied workbook does not preserve the literal Scopus search-history string, extraction timestamp, or a complete screening log. To avoid creating a false audit trail, this article does not present a reconstructed query as though it were the verbatim executed syntax. For submission, the exact Scopus search-history export should be appended to the methods or supplementary material if available. This limitation does not affect the internal accuracy of the descriptive and science-mapping results reported from the verified corpus, but it does constrain full search-level reproducibility.

The analysis proceeds in four stages. First, performance analysis describes annual scientific production, annual citation patterns, leading sources, leading authors, affiliations, country-affiliation frequency, internationally collaborative links, and globally cited documents. Annual citations are interpreted with attention to citation-window effects:

recent publications have had less time to accumulate citations and should not be compared mechanically with older work. Similarly, the 2026 publication count is treated as a partial-year observation.

Second, trend-topic analysis is used to identify the temporal movement of recurrent terms. Biblioshiny reports each term's frequency and the first quartile, median, and third quartile of publication years in which it appears. These distributions make it possible to distinguish earlier infrastructure terms from more recent analytic fronts. Third, keyword co-occurrence analysis maps conceptual proximity. Network prominence is examined through betweenness, closeness, and PageRank measures. The highest PageRank values in the supplied co-word network belong to learning systems, crime, fraud detection, machine learning, and machine-learning, indicating that learning-based fraud analytics occupies the network core.

Fourth, a strategic thematic map is interpreted using Callon centrality and density (Callon et al., 1991; Cobo et al., 2011). Centrality captures the degree to which a thematic cluster connects with other clusters and therefore its importance to the overall field; density represents internal development and cohesion. In the supplied Biblioshiny configuration, 250 merged keywords were analyzed with a minimum frequency of two and Louvain clustering. Four principal clusters emerge: learning systems; artificial intelligence; graph neural networks; and behavioral research. Because lower-left themes can represent either emerging or declining topics, quadrant position alone is insufficient for temporal interpretation. This study therefore triangulates thematic-map location with trend-topic medians and quartiles. Graph neural networks and behavioral research have medians concentrated in 2026 and are thus interpreted as emerging fronts rather than legacy themes in decline.

Indicator	Value
Timespan	2004–2026
Sources	402
Documents	545
Annual growth rate	26.29%

Average citations/document	9.415
References	55,452
Authors	1,825
Author keywords	1,727
Keywords Plus	3,136
Co-authors/document	3.78
International co-authorship	25.32%
Conference papers	296 (54.3%)
Articles	158 (29.0%)
Book chapters	51 (9.4%)
Reviews	23 (4.2%)

Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Table 1. Bibliometric profile of the analyzed corpus

4. Results


4.1 Growth of the Literature and Citation Dynamics

Scientific output is highly uneven over the study period. The corpus begins with one publication in 2004 and one in 2005, remains sparse through 2017, and begins a visible expansion from 2018. Output increases from 9 documents in 2018 to 13 in both 2019 and 2020, 16 in 2021, 28 in 2022, and 30 in 2023. The field then enters a rapid growth phase: 70 documents were published in 2024 and 179 in 2025. The 2025 total is nearly six times the 2023 output. Biblioshiny reports an annual growth rate of 26.29% for the corpus. The 170 documents recorded for 2026 should not be interpreted as an annual decline because the year is incomplete in the supplied dataset.

Citation behavior reinforces the importance of temporal normalization. Older cohorts have had more time to accumulate impact. The 2014 cohort reports a mean of 109 citations per article based on only two documents, while 2018 records 59.11 across nine

documents and 2021 records 46.75 across 16. By contrast, the 2025 and partial-2026 cohorts show 3.58 and 0.17 mean citations per article, respectively. These lower values are expected for recent publications and do not indicate weaker scientific influence. For a fast-growing field, raw citation totals should therefore be read together with publication age, normalized citations, and thematic recency.

The source structure also indicates rapid technical development. Lecture Notes in Networks and Systems is the most productive source with 21 documents, followed by the Lecture Notes in Computer Science family with 16, IEEE Access with 10, IEEE Internet of Things Journal with 7, and Procedia Computer Science with 7. Several conference series and edited proceedings appear among the leading outlets. This pattern is consistent with the document-type profile: conference papers form the majority of the corpus. The field is not yet concentrated in a small set of finance journals; rather, it is distributed across computer science, networking, IoT, cybersecurity, information systems, and finance outlets.



Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Figure 1. Annual scientific production in AI and real-time payment research

4.2 Sources, Authors, Institutions, and Countries

Authorship is similarly dispersed. The most productive authors, Kumar A and Li X, each contribute six documents, while Gupta S, Kumar S, Sandkuhl K, Wang Y, and Zhang Z each contribute five. No author accounts for more than a small fraction of the corpus, suggesting that the domain is being assembled by multiple research communities rather than led by one dominant school. Institutional productivity is also distributed. Excluding records without a reported affiliation, the leading institutions include SRM Institute of Science and Technology (7 documents), IEC College of Engineering and Technology (6), University of Patras (6), University of the Cumberlands (6), and several institutions with five documents, including Beijing University of Posts and Telecommunications, King Saud University, Qatar University, Southeast University, Tsinghua University, and the University of Chinese Academy of Sciences.

Country-affiliation frequency places India first (437 occurrences), followed by the United States (210) and China (177). The United Kingdom (39) and Saudi Arabia (32) form a second tier, while Iraq, Uzbekistan, South Korea, Morocco, Pakistan, Bangladesh, the United Arab Emirates, Australia, Nigeria, Canada, Greece, Iran, Malaysia, and Spain also contribute. These are affiliation-frequency counts rather than mutually exclusive document totals, because internationally co-authored papers contribute to more than one country. International co-authorship is 25.32% overall. The strongest bilateral link is India–United States (22 collaborations), followed by United States–China and India–Saudi Arabia (8 each). Other visible links include United States–Bangladesh and India–Uzbekistan (5 each), as well as several four-document links. The collaboration structure demonstrates a globally distributed field but also suggests that cross-border research could deepen further, particularly because real-time payment infrastructures and financial-crime networks are inherently transnational.


Rank	Source	Documents
1	Lecture Notes in Networks and Systems	21
2	Lecture Notes in Computer Science family	16
3	IEEE Access	10

4	IEEE Internet of Things Journal	7
5	Procedia Computer Science	7
6	Communications in Computer and Information Science	6
7	Lecture Notes in Computer Science	5
8	ACM International Conference Proceeding Series	4
9	Journal of Risk and Financial Management	4
10	ICSES 2024 proceedings	4

Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Table 2. Most productive publication sources

Country Scientific Production



Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Figure 2. Country scientific-production frequency

4.3 Intellectual Influence: Globally Cited Documents

The list of globally cited documents shows that the intellectual base of the corpus is broader than narrow retail-payment fraud. The most cited record is Liu et al. (2019), which applies deep reinforcement learning to performance optimization in blockchain-enabled industrial IoT systems and records 385 citations in the Biblioshiny dataset. Roy et al. (2018), a deep-learning study of credit-card fraud detection, ranks second with 268 citations. Allen et al. (2021), a broad survey of fintech research and policy, ranks third with 252 citations. Other highly cited documents address

neural-network forecasting, blockchain data markets, federated learning, secure IoT blockchain architectures, and the intersection of blockchain and machine learning. This mix indicates that the AI-RTPS field draws intellectual resources from adjacent technology domains: blockchain scalability, distributed computation, IoT, cybersecurity, and financial analytics.

The presence of highly cited adjacent-domain studies is analytically important. It suggests that real-time payment research has not evolved as a closed payment-science specialty. Instead, it imports methods and architectures from wider digital-

infrastructure research and then adapts them to payment requirements. This is particularly visible in reinforcement learning, federated learning, blockchain, and GNNs. The pattern also helps explain why the keyword landscape includes both payment-

specific terms such as instant payment and real-time transactions and technology-general terms such as learning systems, blockchain, network security, and graph neural networks.

Rank	Reference	Document focus	TC	TC/year
1	Liu et al. (2019)	Performance optimization for blockchain-enabled IIoT systems	385	48.13
2	Roy et al. (2018)	Deep learning detecting fraud in credit card transactions	268	29.78
3	Allen et al. (2021)	A survey of fintech research and policy discussion	252	42.00
4	Hernández et al. (2014)	Artificial neural networks for short-term load forecasting	213	16.38
5	Cai et al. (2023)	GTxChain: Secure IoT smart blockchain architecture using GNN	139	34.75
6	Özyilmaz et al. (2018)	IdMob: IoT data marketplace on blockchain	124	13.78
7	Fan et al. (2021)	Hybrid blockchain-based resource trading for federated learning	120	20.00
8	Kayikci & Khoshgoftaar (2024)	Blockchain meets machine learning: A survey	118	39.33
9	Lăzăroiu et al. (2023)	AI/cloud computing in blockchain-based fintech management	115	28.75
10	Yun et al. (2021)	DQN-based optimization for secure sharded blockchain	113	18.83

Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Table 3. Ten most globally cited documents in the Biblioshiny corpus

4.4 Trend Topics, Co-Word Structure, and Thematic Map

Trend-topic analysis provides the clearest view of thematic evolution. “Instant payment” appears as an earlier specialized term, with a median year of 2020. “Internet of things” and “distributed ledger” center around 2022, while “scalability,” “bitcoin,” and “energy utilization” cluster around 2023. “Blockchain,” “intelligent systems,” and “security” become prominent around 2024. The strongest shift


occurs in 2025, when “learning systems” (frequency 176), “crime” (175), and “fraud detection” (152) become dominant, alongside machine-learning and anomaly-detection vocabulary. In 2026, “graph neural networks” (51), “behavioral research” (32), and “financial fraud detections” (22) emerge at the frontier. The temporal sequence shows a move from payment/infrastructure questions toward intelligence, security, and networked fraud analytics.

The co-word network confirms this transition. The learning-systems cluster occupies the central core and contains learning systems (176), crime (175), fraud detection (152), machine learning (144), machine-learning (116), deep learning (90), finance (82), anomaly detection (80), electronic money (74), real-time (68), learning algorithms (46), real-time transactions (41), and financial transactions (37). Learning systems has the highest PageRank in the reported network (0.0693), followed by crime (0.0661) and fraud detection (0.0533). Machine learning and its hyphenated variant also rank highly. These metrics indicate that the conceptual center of the field is now the use of learning methods to detect suspicious or anomalous financial activity under real-time conditions.

A second major cluster is organized around artificial intelligence, blockchain, and security. Its leading terms include artificial intelligence (149), blockchain (146), block-chain (112), network security (92), data privacy (35), cyber security (32), IoT (30), cybersecurity (28), scalability (27), security (25), decentralized finance (25), federated learning (24),


economic and social effects (23), decentralization (23), and regulatory compliance (22). The combination indicates that AI in real-time payments is inseparable from the architecture on which data and decisions are distributed. Security, privacy, compliance, and scale are not peripheral constraints; they are part of the thematic core.

The strategic thematic map condenses the field into four clusters. The learning-systems cluster has the highest Callon centrality (11.06) and density (28.72), making it the strongest motor theme. The artificial-intelligence cluster is also a motor theme with centrality 9.60 and density 27.17. Behavioral research has moderate centrality (6.99) and density (25.51), while GNNs have lower centrality (3.51) and density (22.04). On quadrant position alone, the latter two could be classified as emerging or declining. Their temporal evidence resolves the ambiguity: both are recent, with trend-topic medians concentrated in 2026. They should therefore be interpreted as emerging themes that have not yet developed the internal density or cross-field centrality of the two mature motor themes.



Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Figure 3. Trend topics and their temporal concentration




Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Figure 4. Keyword co-occurrence network

Cluster	Centrality	Density	Freq.	Representative terms	Interpretation
Learning systems	11.06	28.72	2,156	learning systems; crime; fraud detection; machine learning; deep learning; anomaly detection	Mature motor theme
Artificial intelligence	9.60	27.17	1,426	AI; blockchain; network security; data privacy; cybersecurity; scalability; federated learning	Mature motor theme
Behavioral research	6.99	25.51	589	behavioral research; banking; big data; decision making; fintech; risk management	Emerging (recent temporal profile)
Graph neural networks	3.51	22.04	280	GNN; reinforcement learning; graph neural network; benchmarking; phishing; heterogeneous graph	Emerging frontier

Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Table 4. Strategic thematic clusters and interpretation



Source: Authors' analysis of the supplied Scopus-derived Biblioshiny report.

Figure 5. Strategic thematic map of AI in real-time payment research

5. DISCUSSION

The bibliometric evidence points to a field undergoing a structural transition. Early research was distributed among instant payments, computers, distributed ledgers, IoT, and blockchain-oriented infrastructure. The current core is different: learning systems, crime, fraud detection, anomaly detection, security, and real-time transactions now dominate. This transition is consistent with the operational maturation of fast-payment rails. Once the basic infrastructure for near-instant transfer becomes available, the research problem shifts from whether and how to move value quickly to how to preserve trust, detect abuse, and manage risk at the same speed.

The first major implication is that AI for RTPS is becoming a real-time decision problem rather than a conventional predictive-model problem. In many financial ML studies, models are evaluated off-line on static datasets and judged by accuracy, F1 score, area under the curve, or related metrics. Those measures are necessary but insufficient for live payment systems. Ryman-Tubb et al. (2018) emphasized the gap between academic fraud-detection performance and operational deployment, including the importance of real-time functionality and business-oriented metrics. The present thematic results show that this

concern has become more central, not less. “Real-time,” “real-time transactions,” anomaly detection, crime, and learning systems now co-occur in the network core. Future studies should therefore report end-to-end inference latency, throughput, alert volume, false-positive cost, time-to-detect, calibration, and resilience under drift in addition to predictive accuracy.

Second, the rise of GNNs signals a move from isolated transaction scoring toward relational intelligence. Traditional fraud models often assume that each transaction is an independent observation with a feature vector. Real financial crime violates that assumption. Mule accounts, synthetic identities, coordinated merchant abuse, account takeovers, layering, and scam networks produce patterns across entities and time. GNNs can encode relationships among accounts, devices, beneficiaries, merchants, IP addresses, or institutions. Motie and Raahemi (2024) show that graph-based financial-fraud research is promising but still underdeveloped in unsupervised, dynamic, and multi-level settings. The Biblioshiny trend data place GNNs precisely where one would expect an emerging frontier: recent, visible, but not yet as dense or central as conventional learning-based fraud detection.

Third, the behavioral-research cluster broadens the unit of analysis from transactions and networks to people. This is increasingly important because many contemporary payment losses arise when legitimate customers are manipulated into authorizing transfers. In such cases, authentication may work exactly as designed while the underlying payment intent has been socially engineered. The 2025 joint EBA–ECB fraud assessment reports that strong customer authentication remains effective against the fraud types it was designed to address, while payer manipulation is becoming more important (European Banking Authority & European Central Bank, 2025). Behavioral AI can complement transaction analytics by modeling changes in user interaction, beneficiary novelty, payment context, conversational cues, or deviations from a customer’s normal decision process. The challenge is to do so without producing intrusive surveillance or discriminatory outcomes.

Fourth, privacy-preserving collaboration is likely to become a core architectural issue. Individual payment service providers see only part of a criminal network, while payment-system operators may observe a wider topology. The artificial-intelligence cluster already includes data privacy, federated learning, cybersecurity, and regulatory compliance. Project Hertha demonstrates the value of system-level analytics but also highlights legal, practical, and governance constraints (BIS Innovation Hub, 2025). A promising research direction is therefore federated or distributed learning in which institutions improve models collaboratively without centralizing raw customer data. Such systems will need robust defenses against data poisoning, model inversion, membership inference, and uneven data quality across participants.

Fifth, explainability should be treated as an operational control rather than an after-the-fact visualization. In RTPS, a model decision may block a payment, trigger step-up authentication, delay settlement, prioritize an alert, or influence a compliance investigation. Each action can affect customers and institutions materially. Research on XAI in finance emphasizes the growing importance of interpretability, trust, and risk assessment (Chen et al., 2023). For real-time payments, explainability must also be time-sensitive: explanations should be available quickly enough to support operational

decisions and sufficiently stable for post-event audit. Future model comparisons should therefore measure explanation fidelity, stability, actionability, and computational overhead alongside predictive performance.

Sixth, the corpus reveals convergence between fraud intelligence and payment-system optimization. Reinforcement learning and intelligent-agent terms appear alongside blockchain, scalability, and real-time systems. Liu et al. (2019), the corpus’s most globally cited document, demonstrates the influence of deep reinforcement learning in optimizing complex distributed environments. More recently, Aldasoro and Desai (2025) test a generative AI agent for cash management in payment systems, showing that the research frontier is extending to liquidity and operational decision support. This creates a new governance frontier. A fraud model typically recommends whether a transaction is suspicious; an agentic system may recommend how to allocate liquidity, prioritize queues, or respond to operational constraints. The latter has a larger action space and potentially broader systemic consequences. Human oversight, constrained action policies, scenario testing, and rollback mechanisms should therefore be designed into AI-enabled payment operations from the outset.

Finally, the geographic results point to both strength and imbalance. India, the United States, and China dominate affiliation frequency, and the India–United States link is the largest collaborative dyad. This is understandable given the scale of digital-payment adoption and technology research in these countries, but generalizability requires wider evidence. Payment behavior, fraud typologies, identity infrastructures, regulation, connectivity, and consumer protection differ across jurisdictions. Cross-country comparative research, particularly involving emerging markets with rapidly expanding instant-payment systems, can test whether AI models and governance practices transfer across institutional contexts.

6. FUTURE RESEARCH AGENDA

The combined thematic and temporal evidence supports a research agenda organized around seven connected fronts. The objective should not be simply to build more accurate classifiers, but to create AI systems that remain useful under real-time latency,

adaptive adversaries, networked crime, privacy constraints, and institutional accountability. Table 5 translates the bibliometric findings into testable research directions.

Front	Unresolved problem	Recommended approach	Expected contribution
Temporal intelligence	Concept drift, delayed labels, sequence dependence	Streaming/online learning; temporal validation; drift detectors	Lower detection latency with stable performance under evolving fraud
Graph/network intelligence	Institution-level models miss coordinated networks	Temporal/heterogeneous GNNs; self-supervised graph learning	Detection of mule, scam and laundering networks across linked entities
Privacy-preserving collaboration	Raw cross-institution data cannot be freely pooled	Federated learning; secure aggregation; MPC; differential privacy	Network-level intelligence with controlled disclosure
Explainable real-time decisioning	Black-box actions are difficult to audit or contest	Interpretable models; fast post-hoc explanations; counterfactuals	Actionable explanations with measured fidelity, stability and latency
Behavioral scam detection	Legitimate users can be manipulated into authorizing fraud	Behavioral sequence models; contextual risk; human-in-the-loop design	Earlier identification of anomalous intent without excessive friction
Adversarial resilience	Attackers adapt to models and may poison data	Adversarial testing; red teaming; robust learning; feedback-loop controls	Resilience to evasion, poisoning and synthetic fraud patterns
System/agent intelligence	AI is moving into liquidity, routing and operations	Constrained RL/LLM agents; digital twins; scenario testing	Efficiency gains with bounded action spaces, override and rollback

Source: Authors' synthesis based on the bibliometric and thematic findings and the cited literature.

Table 5. Research agenda for AI-enabled real-time payment systems

7. Research, Managerial, and Policy Implications

For researchers, the main implication is methodological. Static benchmark datasets and random train–test splits are increasingly misaligned with the problem structure. RTPS studies should use chronological validation, delayed-label simulation, concept-drift analysis, cost-sensitive evaluation, and, where relevant, temporal or heterogeneous graphs. Claims of “real-time” capability should be supported by measured latency and throughput under realistic load rather than inferred from model type. Synthetic data can be valuable when real payment data cannot be shared, but realism, privacy risk, and distributional fidelity should be documented.

For payment service providers and system operators, the findings favor layered controls rather than a single monolithic model. High-speed rules and lightweight models can perform first-pass screening; specialized anomaly or graph models can examine higher-risk events; human investigators can focus on the smallest, most consequential set of alerts. System-level analytics may complement institution-level models by exposing cross-provider relationships. However, governance must specify who can act on system-wide signals, how false positives are handled, and how customer rights and confidentiality are protected.

For regulators and central banks, the emerging themes suggest that AI governance in payments will increasingly intersect with operational resilience,

model risk, data protection, and financial-crime supervision. The relevant question is not whether AI is used, but where it sits in the decision chain and what happens when it fails. Supervisory frameworks should distinguish advisory models from models that directly influence authorization, settlement, liquidity, or customer access. Documentation of data lineage, performance drift, human override, incident response, and explainability will become increasingly important as AI moves closer to the payment core.

8. Limitations and Reproducibility Note

This study has four limitations. First, it relies on a Scopus-derived corpus; records unique to Web of Science, IEEE Xplore, ACM Digital Library, SSRN, working-paper repositories, or industry sources may be underrepresented. Second, the supplied Biblioshiny report does not retain the verbatim Scopus query, extraction timestamp, or full screening log. The analysis therefore reports only verified corpus-level outputs and avoids reconstructing unverified audit details. Third, 2026 is an incomplete publication year, and recent citation counts are affected by short exposure time. Fourth, co-word and thematic maps depend on author/indexer terminology. Variants such as “machine learning” and “machine-learning,” or “blockchain” and “block-chain,” can fragment frequencies despite keyword merging. Future replications should preserve the full search history, document preprocessing rules explicitly, normalize lexical variants before mapping, and compare Scopus results with at least one complementary database.

CONCLUSION

Research on artificial intelligence in real-time payment systems has moved from a small, infrastructure-oriented literature into a rapidly expanding interdisciplinary field. The 545-document corpus shows an inflection after 2023, with output rising sharply in 2024 and 2025. Its current conceptual core is organized around learning systems, crime, fraud detection, AI, blockchain, and security. These are mature motor themes with high thematic centrality and density. Graph neural networks and behavioral research are less central but temporally newer, making them credible emerging fronts rather than declining topics.

The field’s next phase will be defined by a shift from isolated transaction classification to adaptive payment intelligence. Such intelligence will need to understand temporal sequences, relational networks, user behavior, and operational context while meeting strict requirements for latency, privacy, explainability, and resilience. The most consequential research opportunities therefore lie at the boundaries between technical performance and payment-system governance: dynamic GNNs for coordinated fraud, federated analytics across institutions, explainable low-latency decisioning, behavioral models for manipulated payments, adversarially robust learning, and carefully constrained AI agents for liquidity and operational management. A mature AI–RTPS research program will be judged not only by whether a model detects more anomalies, but by whether it improves the safety, efficiency, accountability, and trustworthiness of real-time financial infrastructures.

Data availability statement. The bibliometric results reported in this article were derived from the Scopus/Biblioshiny workbook supplied for the study. The exact Scopus search-history export should be retained with the submission materials to support full replication.

Conflict of interest. The authors declare no conflict of interest.

Funding. No external funding is reported in the supplied study materials.

REFERENCES

1. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113. <https://doi.org/10.1016/j.jnca.2016.04.007>
2. Aldasoro, I., & Desai, A. (2025). AI agents for cash management in payment systems (BIS Working Papers No. 1310). Bank for International Settlements.
3. Allen, F., Gu, X., & Jagtiani, J. (2021). A survey of fintech research and policy discussion. *Review of Corporate Finance*, 1(3–4), 259–339. <https://doi.org/10.1561/114.000000007>
4. Aria, M., & Cuccurullo, C. (2017). bibliometrix: An R-tool for comprehensive science mapping analysis. *Journal of Informetrics*, 11(4), 959–975. <https://doi.org/10.1016/j.joi.2017.08.007>

5. Aria, M., Cuccurullo, C., D'Aniello, L., & Spano, M. (2026). Biblioshiny and the SAAS workflow: An integrated framework for transparent and reproducible science mapping—A demonstration through the replication of a study. *Journal of Informetrics*.
<https://doi.org/10.1016/j.joi.2026.101837>
6. Bank for International Settlements Innovation Hub. (2025). Project Hertha: Identifying financial crime patterns in real-time retail payment systems. Bank for International Settlements.
7. Brogi, M., & Lagasio, V. (2024). New but naughty: The evolution of misconduct in FinTech. *International Review of Financial Analysis*, 95, 103489.
<https://doi.org/10.1016/j.irfa.2024.103489>
8. Cai, J., Liang, W., Li, X., Li, K.-C., Gui, Z., & Khan, M. K. (2023). GTxChain: A secure IoT smart blockchain architecture based on graph neural network. *IEEE Internet of Things Journal*, 10(24), 21502–21514.
<https://doi.org/10.1109/JIOT.2023.3296469>
9. Callon, M., Courtial, J.-P., & Laville, F. (1991). Co-word analysis as a tool for describing the network of interactions between basic and technological research: The case of polymer chemistry. *Scientometrics*, 22, 155–205.
<https://doi.org/10.1007/BF02019280>
10. Chen, X.-Q., Ma, C.-Q., Ren, Y.-S., Lei, Y.-T., Huynh, N. Q. A., & Narayan, S. (2023). Explainable artificial intelligence in finance: A bibliometric review. *Finance Research Letters*, 56, 104145.
<https://doi.org/10.1016/j.frl.2023.104145>
11. Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). An approach for detecting, quantifying, and visualizing the evolution of a research field: A practical application to the fuzzy sets theory field. *Journal of Informetrics*, 5(1), 146–166.
<https://doi.org/10.1016/j.joi.2010.10.002>
12. Committee on Payments and Market Infrastructures. (2016). Fast payments—Enhancing the speed and availability of retail payments (CPMI Papers No. 154). Bank for International Settlements.
13. Committee on Payments and Market Infrastructures. (2021). Developments in retail fast payments and implications for RTGS systems. Bank for International Settlements.
14. Cornelli, G., Frost, J., Warren, J., Yang, C., & Velásquez, C. (2024). Retail fast payment systems as a catalyst for digital finance (BIS Working Papers No. 1228; revised June 2025). Bank for International Settlements.
15. Desai, A., Kosse, A., & Sharples, J. (2024). Finding a needle in a haystack: A machine learning framework for anomaly detection in payment systems (BIS Working Papers No. 1188). Bank for International Settlements.
16. Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296.
<https://doi.org/10.1016/j.jbusres.2021.04.070>
17. European Banking Authority, & European Central Bank. (2025). 2025 report on payment fraud. European Banking Authority and European Central Bank.
18. Fan, S., Zhang, H., Zeng, Y., & Cai, W. (2021). Hybrid blockchain-based resource trading system for federated learning in edge computing. *IEEE Internet of Things Journal*, 8(4), 2252–2264.
<https://doi.org/10.1109/JIOT.2020.3028101>
19. Frost, J., Koo Wilkens, P., Kosse, A., Shreeti, V., & Velásquez, C. (2024). Fast payments: Design and adoption. *BIS Quarterly Review*, March.
20. Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577.
<https://doi.org/10.1016/j.jbef.2021.100577>
21. Hernández, L., Baladrón, C., Aguiar, J. M., Carro, B., Sánchez-Esguevillas, A., & Lloret, J. (2014). Artificial neural networks for short-term load forecasting in microgrids environment. *Energy*, 75, 252–264.
<https://doi.org/10.1016/j.energy.2014.07.065>
22. Kayikci, S., & Khoshgoftaar, T. M. (2024). Blockchain meets machine learning: A survey. *Journal of Big Data*, 11, 9.
<https://doi.org/10.1186/s40537-023-00852-y>
23. Lăzăroiu, G., Bogdan, M., Geamănu, M., Hurloiu, L., Ionescu, L., & Ștefănescu, R. (2023).

- Artificial intelligence algorithms and cloud computing technologies in blockchain-based fintech management. *Oeconomia Copernicana*, 14(3), 707–730. <https://doi.org/10.24136/oc.2023.021>
24. Liu, M., Yu, F. R., Teng, Y., Leung, V. C. M., & Song, M. (2019). Performance optimization for blockchain-enabled industrial Internet of Things (IIoT) systems: A deep reinforcement learning approach. *IEEE Transactions on Industrial Informatics*, 15(6), 3559–3570. <https://doi.org/10.1109/TII.2019.2897805>
 25. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, 122156. <https://doi.org/10.1016/j.eswa.2023.122156>
 26. Özyılmaz, K. R., Doğan, M., & Yurdakul, A. (2018). IDMoB: IoT data marketplace on blockchain. In 2018 Crypto Valley Conference on Blockchain Technology (CVCBT) (pp. 11–19). IEEE. <https://doi.org/10.1109/CVCBT.2018.00007>
 27. Pattnaik, D., Ray, S., & Raman, R. (2024). Applications of artificial intelligence and machine learning in the financial services industry: A bibliometric review. *Heliyon*, 10(1), e23492. <https://doi.org/10.1016/j.heliyon.2023.e23492>
 28. Roy, A., Sun, J., Mahoney, R., Alonzi, L., Adams, S., & Beling, P. (2018). Deep learning detecting fraud in credit card transactions. In 2018 Systems and Information Engineering Design Symposium (SIEDS) (pp. 129–134). IEEE. <https://doi.org/10.1109/SIEDS.2018.8374722>
 29. Ryman-Tubb, N. F., Krause, P., & Garn, W. (2018). How artificial intelligence and machine learning research impacts payment card fraud detection: A survey and industry benchmark. *Engineering Applications of Artificial Intelligence*, 76, 130–157. <https://doi.org/10.1016/j.engappai.2018.07.008>
 30. Yun, J., Goh, Y., & Chung, J.-M. (2021). DQN-based optimization framework for secure sharded blockchain systems. *IEEE Internet of Things Journal*, 8(2), 708–722. <https://doi.org/10.1109/JIOT.2020.3006896>
 31. Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429–472. <https://doi.org/10.1177/1094428114562629>

HOW TO CITE: Vijaya Kumar*, Artificial Intelligence In Real-Time Payment Systems: A Bibliometric And Thematic Analysis Of Emerging Research Trends (2004–2026), *Int. J. Sci. R. Tech.*, 2026, 3 (8), 739-754. <https://doi.org/10.5281/zenodo.21991392>