

Artificial Intelligence in Real-Time Payment Systems: A Bibliometric and Thematic Analysis of Emerging Research Trends

Swathi N ¹, Bora Upendra Rao ²

¹Research Scholar, School of Commerce, Presidency University, Bangalore, Karnataka, India

²Assistant Professor, School of Commerce, Presidency University, Bangalore, Karnataka, India

Abstract:

The rapid diffusion of real-time payment systems (RTPs) has compressed payment authorization, clearing, and fraud-control decisions into near-instantaneous processes, increasing demand for artificial intelligence (AI)-enabled risk detection capable of operating under severe latency, scalability, and security constraints. Although research on digital payments, financial fraud, blockchain, and AI has expanded considerably, the intellectual structure and thematic development of scholarship specifically connecting AI with real-time and instant payment environments remain insufficiently mapped. This study conducts a bibliometric and thematic analysis of 545 documents published between 2004 and 2026, using outputs generated through the Bibliometrix/Biblioshiny environment in R. Performance analysis, source and author productivity, country and institutional contributions, global citation analysis, keyword frequency, trend-topic analysis, co-word mapping, thematic mapping, bibliographic coupling, and international collaboration networks are examined. The corpus records an annual growth rate of 26.29%, while 76.9% of all publications appeared during 2024–2026, indicating a sharply accelerating research domain. Conference papers account for 54.3% of the corpus, reflecting a technology-intensive and rapidly evolving knowledge base. Thematic mapping identifies two principal motor themes: learning systems–fraud detection–machine learning and artificial intelligence–blockchain–security. Graph neural networks and reinforcement learning form a lower-centrality but highly recent research front, while behavioral research, FinTech analytics, risk management, and decision-making constitute a second emerging trajectory. Trend analysis reveals a temporal transition from early interest in instant payments toward blockchain, distributed infrastructure, scalability, and security, followed by a pronounced 2025–2026 shift toward fraud detection, learning systems, and graph neural networks. Bibliographic coupling further identifies three active fronts: operational fraud/anomaly detection, machine/deep-learning-based detection, and blockchain/security architectures. The study concludes that AI research in RTPs is moving from transaction-level classification toward relational, adaptive, privacy-aware, explainable, and operationally deployable intelligence. A future research agenda is proposed around temporal graph learning, concept drift, adversarial robustness, federated learning, explainability, latency-aware benchmarking, cross-institutional datasets, and responsible AI governance.

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

1. Introduction

Payment infrastructures are undergoing a fundamental shift from batch-oriented settlement and delayed fund availability toward systems in which payment messages are transmitted, processed, and made available to recipients in real or near-real time. The Committee on Payments and Market Infrastructures (CPMI) defines fast payments as payment arrangements in which funds become available to the payee immediately or almost immediately and services operate on a near-24/7 basis (Committee on Payments and Market Infrastructures [CPMI], 2016). Subsequent developments have expanded the operational significance of these infrastructures, including closer integration with real-time settlement, standardized messaging, broader participation, and increasingly continuous payment availability (CPMI, 2021).

Speed, however, changes the economics and technological conditions of payment risk. In conventional payment environments, institutions may have minutes, hours, or longer to identify anomalous behavior, suspend settlement, obtain supplementary information, or initiate manual review. Real-time payments narrow this decision window

considerably. Fraud controls must increasingly make high-quality decisions before or during authorization rather than relying predominantly on post-transaction intervention. The combination of immediate availability of funds, high transaction volumes, continuously operating infrastructure, diverse customer behaviors, and evolving attack patterns therefore creates a setting in which traditional static rules alone can become insufficient.

Artificial intelligence (AI), encompassing machine learning (ML), deep learning (DL), anomaly detection, graph learning, reinforcement learning, and related computational approaches, has consequently become central to the technological discussion surrounding payment risk. Earlier fraud-detection systems frequently relied on expert rules and conventional classifiers, whereas contemporary systems increasingly combine behavioral signals, transaction histories, network relationships, device information, and adaptive models. Diadiushkin et al. (2019) explicitly identified the short processing time available in instant payments as a distinctive fraud-detection challenge and examined AI as a mechanism for automating fraud discovery under these conditions.

The technical problem is more complex than simply maximizing classification accuracy. Fraud data are typically highly imbalanced; fraudulent strategies evolve; transaction streams exhibit temporal dependence; labels may arrive with delay; false positives impose financial and customer-experience costs; and production systems must operate within strict latency and availability requirements. Data engineering therefore becomes inseparable from model design (Baesens et al., 2021), while anomaly-detection research emphasizes the difficulty of identifying changing and sparsely represented abnormal behavior (Hilal et al., 2022). Reviews of credit-card and financial fraud detection similarly document an evolution from classical supervised classifiers toward deep, ensemble, graph-based, and increasingly adaptive methods (Cherif et al., 2023; Motie & Raahemi, 2024).

The growing relevance of relational information is especially consequential. Fraud is rarely an isolated property of a single transaction. Accounts, merchants, devices, beneficiaries, IP addresses, payment instruments, and counterparties form dynamic networks through which fraudulent behavior can propagate. Graph neural networks (GNNs) provide a mechanism for exploiting these relationships. For example, BRIGHT was developed to enable graph-based real-time fraud inference while preventing future information from leaking into historical predictions and reducing the latency associated with graph querying and inference (Lu et al., 2022). More recent work has extended this direction through heterogeneous temporal graph architectures specifically targeting real-time transaction fraud (Nguyen & Le, 2025).

Despite these developments, the relevant scholarship remains distributed across payment systems, computer science, banking, cybersecurity, blockchain, financial fraud, Internet of Things (IoT), and FinTech. This fragmentation makes it difficult to determine whether AI in RTPs is developing into a coherent research domain or remains an intersection of adjacent literatures.

Existing bibliometric studies only partly address this problem. Moura et al. (2025) map the broader relationship between AI and financial fraud prevention, whereas Tanwar and Arora (2025) examine the evolution of digital payment systems broadly. The latter analyzed 871 Scopus publications from 2003–2023 using Bibliometrix/Biblioshiny and focused substantially on digital payment adoption and behavioral constructs. These studies provide important foundations, but neither isolates the AI–real-time-payment intersection as the primary unit of analysis.

The present study therefore offers a more tightly delimited mapping of the emerging knowledge space connecting artificial intelligence, fraud analytics, security technologies, and real-time payment environments. Rather than treating digital payments generally or fraud detection independently of payment speed, it investigates how research themes converge around the specific technological requirements of real-time and instant transactions.

The study addresses five research questions:

RQ1. How has research on AI in real-time payment systems evolved in terms of publication growth, document types, sources, authors, institutions, countries, and citations?

RQ2. What conceptual structures dominate the field according to keyword frequencies and co-word relationships?

RQ3. How have research topics evolved over time, and which themes represent established versus emerging areas?

RQ4. What major contemporary research fronts are revealed by bibliographic coupling and international collaboration patterns?

RQ5. What theoretical, technological, methodological, and governance priorities should guide future research on AI-enabled real-time payment systems?

By answering these questions, the study contributes in four ways. First, it provides a quantitative performance profile of a rapidly expanding literature. Second, it identifies the field's conceptual architecture through co-word and thematic analysis. Third, it traces a temporal movement from general payment infrastructure toward increasingly sophisticated real-time fraud intelligence. Fourth, it develops a future research framework in which model accuracy is considered alongside latency, concept drift, relational structure, privacy, explainability, adversarial robustness, and regulatory accountability.

2. Conceptual Background and Literature Review

2.1 Real-time payment systems as a distinct technological environment

Real-time or fast payment systems should not be treated merely as faster versions of conventional electronic payments. Their continuous availability and near-immediate fund transfer alter operational requirements across clearing, settlement, liquidity management, security, and fraud prevention. CPMI (2016) emphasizes immediate fund availability and continuous accessibility as defining characteristics, while later CPMI work documents the increasing interaction between retail fast-payment infrastructures and real-time settlement arrangements (CPMI, 2021). These characteristics produce a distinct risk architecture. A false negative may allow funds to leave the victim's effective control almost immediately, whereas an excessive false-positive rate can undermine customer trust in a payment channel whose value proposition is speed and convenience. Fraud controls must therefore optimize multiple objectives simultaneously: detection accuracy, precision, recall, latency, throughput, interpretability, customer friction, and financial loss.

This context also changes the role of historical information. Fraud models must distinguish legitimate behavioral novelty from malicious novelty. A customer making an unusual but valid instant payment may resemble a fraudulent account takeover, while sophisticated fraudsters may deliberately imitate legitimate patterns. Accordingly, RTP fraud detection is inherently a streaming, cost-sensitive, dynamically changing classification problem rather than a static binary classification exercise.

2.2 Artificial intelligence and fraud analytics

Machine learning offers several advantages over fixed-rule systems. Supervised algorithms can identify nonlinear associations across large numbers of transactional features; unsupervised and semi-supervised methods can detect unusual behavior when labeled fraud examples are limited; ensemble approaches can combine complementary models; and deep learning can learn complex representations from high-dimensional data.

Yet financial fraud presents several persistent methodological difficulties. Class imbalance can result in misleadingly high overall accuracy even when models detect relatively few fraudulent transactions. Fraud patterns are non-stationary, meaning models trained on historical data may deteriorate as attacker behavior changes. Ground-truth labels can be incomplete or delayed, and the economic consequences of false positives and false negatives are unequal. Hilal et al. (2022) consequently position anomaly detection as an important component of financial-fraud research, while Cherif et al. (2023) show how disruptive technologies have broadened the fraud-detection toolkit beyond conventional classifiers.

Data preparation is equally consequential. Baesens et al. (2021) argue that data engineering decisions materially determine fraud-detection performance because raw transactional records must be transformed into representations capable of capturing behavioral context. This becomes particularly important for real-time systems, where features must be produced quickly enough to support operational authorization rather than retrospective analysis.

The historical progression is visible in the broader literature. Deep-learning approaches were already demonstrating the potential of neural architectures for credit-card fraud detection by the late 2010s (Roy et al., 2018). Current research increasingly focuses on systems designed for continuous transaction streams and deployment constraints rather than offline model comparison alone.

2.3 From independent transactions to relational intelligence

An important development involves conceptualizing fraud as a network phenomenon. Transaction-level tabular models generally treat observations as independent after feature generation. Financial ecosystems, however, are relational. A single device may connect multiple accounts; beneficiaries may receive money from apparently unrelated senders; fraudulent merchant networks may share infrastructure; and money-mule accounts may form temporally evolving transaction communities. Graph neural networks can encode these relationships directly. Tong and Shen (2023), for example, combine graph learning with imbalance-sensitive fraud detection, while the systematic review by Motie and Raahemi (2024) demonstrates the increasing prominence of GNN approaches across financial fraud applications.

The real-time environment introduces an additional constraint: graph models cannot simply aggregate arbitrary neighboring information because such information may not have existed at the moment a payment decision was required. Lu et al.'s (2022) BRIGHT architecture responds to this challenge by separating historical graph representation from real-time transaction inference. The study reports substantial reductions in high-percentile inference latency while maintaining predictive advantages over traditional alternatives. Its importance is conceptual as much as computational: it demonstrates that graph-based fraud models must be designed around temporal causality and production latency, not merely graph predictive accuracy.

More recently, Nguyen and Le (2025) develop a heterogeneous temporal GNN for real-time transaction fraud detection. The appearance of these studies among the strongly coupled recent documents in the present corpus is consistent with a broader movement from isolated transaction scoring toward temporal-network intelligence.

2.4 Blockchain, security, and distributed payment architectures

A second stream connecting AI and payments emerges through blockchain, distributed ledgers, network security, decentralized finance, authentication, privacy, and smart contracts. Blockchain is not synonymous with real-time retail payment infrastructure, but bibliometrically the concepts intersect because both concern digital value transfer, distributed trust, automated verification, and new financial architectures. AI within this stream plays several roles. It supports intrusion detection, anomaly recognition, identity verification, transaction monitoring, risk scoring, and predictive cybersecurity. Conversely, distributed technologies are sometimes proposed as mechanisms for data integrity, traceability, decentralized trust, and collaborative financial services. The intersection also introduces tensions. Distributed architectures can improve traceability while simultaneously creating privacy issues; strong identity controls may reduce fraud but increase surveillance risk; decentralized systems can remove centralized bottlenecks while complicating governance and accountability. The importance of blockchain, network security, data privacy, cybersecurity, regulatory compliance, authentication, and smart contracts in the thematic results of this study indicates that the AI–RTP literature is not limited to fraud classification but includes the architecture within which intelligent payment controls operate.

2.5 Need for a dedicated bibliometric synthesis

Bibliometric analysis is particularly appropriate for rapidly expanding and interdisciplinary research domains because it complements conventional narrative review with quantitative evidence about publication structure, intellectual relationships, conceptual clusters, and thematic evolution. Donthu et al. (2021) distinguish performance analysis from science mapping, while Zupic and Čater (2015) demonstrate how bibliometric networks can reveal intellectual, conceptual, and social structures not readily visible through narrative review alone. Bibliometrix provides an integrated R-based framework for implementing such analyses (Aria & Cuccurullo, 2017). The need is particularly strong here because prior scholarship divides across several literatures. Research on digital payment systems broadly includes adoption, consumer behavior, mobile payments, financial inclusion, infrastructure, and trust. Research on AI-enabled financial fraud includes banking, insurance, credit

cards, anti-money laundering, accounting fraud, e-commerce, and cybersecurity. The specific intersection between AI and real-time payment processing is therefore easily obscured within broader reviews.

The present analysis addresses that problem by examining the bibliometric structure of the search-defined corpus and identifying which parts of the larger FinTech, cybersecurity, and AI landscape are actually converging around real-time transaction processing.

3. Methodology

3.1 Research design

The study adopts a bibliometric science-mapping design combining performance analysis with conceptual and relational mapping. This distinction follows established bibliometric guidance: performance analysis evaluates the productivity and influence of scientific actors, whereas science mapping investigates relationships among authors, documents, references, keywords, and themes (Donthu et al., 2021; Zupic & Čater, 2015). The analytical outputs were generated in the Bibliometrix/Biblioshiny environment. Bibliometrix is an R-based scientific mapping framework supporting descriptive bibliometrics, citation analysis, collaboration networks, co-word analysis, thematic mapping, and related techniques (Aria & Cuccurullo, 2017). The final analytical corpus supplied for this study contains 545 documents covering 2004–2026 and 402 publication sources.

A publication-quality bibliometric study should report the precise database query, date of retrieval, language/document filters, and screening decisions so that the corpus can be independently reproduced. The supplied Biblioshiny analytical workbook does not retain the original Scopus query syntax or screening log. Accordingly, these elements are not reconstructed retrospectively in this article. Before final journal submission, the exact Scopus search string and extraction date from the original search record should be inserted into the study protocol or supplementary material.

3.2 Corpus characteristics

The Biblioshiny report records 55,452 cited references, 3,136 Keywords Plus/index keywords, 1,727 author keywords, and 1,825 authors. The average number of co-authors per document is 3.78, while international co-authorship accounts for 25.32% of the corpus. Document types are heterogeneous, with conference papers constituting the largest category. This composition is particularly relevant because computer science, cybersecurity, and AI frequently communicate emerging technical developments through conference proceedings before journal consolidation.

Table 1. Bibliometric profile of the corpus

Indicator	Result
Timespan	2004–2026
Documents	545
Sources	402
References	55,452
Annual growth rate	26.29%
Average document age	1.95 years
Average citations per document	9.415
Authors	1,825
Author keywords	1,727

Keywords Plus/index keywords	3,136
Co-authors per document	3.78
International co-authorship	25.32%
Single-authored documents	65
Conference papers	296 (54.3%)
Journal articles	158 (29.0%)
Book chapters	51 (9.4%)
Reviews	23 (4.2%)
Other document types	17 (3.1%)

Source: Authors' analysis of the supplied Biblioshiny report.

3.3 Analytical procedures

The analysis comprised seven complementary components. First, annual scientific production was used to identify the temporal development and acceleration of the field. Second, source, author, affiliation, and country productivity were examined to identify the principal publication channels and knowledge-producing actors. Third, global citation analysis was used to identify highly visible documents. Raw citation counts were interpreted cautiously because older publications have had more time to accumulate citations. Fourth, keyword frequency and trend-topic analysis were employed to trace the conceptual vocabulary of the field and changes in the temporal prominence of specific topics. Fifth, co-word analysis was used to examine conceptual proximity between terms. In co-word mapping, frequent co-occurrence indicates that concepts are repeatedly used together within the same scholarly documents, providing evidence about the conceptual organization of a research domain. Sixth, thematic mapping was interpreted using Callon centrality and Callon density. Centrality indicates the degree to which a theme connects with other themes, whereas density reflects the internal development or cohesiveness of the theme. The resulting strategic diagram distinguishes highly developed and central motor themes, highly developed but isolated niche themes, central but less-developed basic themes, and themes that may be emerging or declining (Cobo et al., 2011). For the supplied thematic analysis, the Biblioshiny settings were: merged keywords (KW_Merged), 250 terms, minimum frequency of two, one-word n-grams, no stemming, and Louvain clustering. Seventh, bibliographic coupling was applied at the document level. Documents sharing references are bibliographically coupled, and dense coupling among recent publications can reveal contemporary research fronts. The supplied analysis used cited references (CR), 250 documents, minimum frequency of two, local citation impact, and Walktrap clustering. Finally, an international collaboration network was examined to identify leading bilateral research relationships.

3.4 Interpretation and data-quality considerations

Bibliometric findings depend on the quality and normalization of metadata. The supplied output retains several orthographic variants, including “machine learning” and “machine-learning,” “blockchain” and “block-chain,” and variants of “real time.” These terms should ideally be consolidated through a thesaurus before final production of figures. Their presence does not invalidate the network structure but can divide frequencies and weaken apparent links.

Several broad index terms—including “crime,” “learning systems,” and some peripheral technical vocabulary—also reflect database indexing rather than exclusively author-selected terminology. Therefore, cluster labels are interpreted together with their constituent keywords rather than literally. Country frequencies are treated as **country-affiliation occurrences**, not mutually exclusive document counts, because internationally co-authored

documents may contribute to more than one country. Similarly, 2026 represents an incomplete publication year and should not be directly compared with the full 2025 calendar year.

4. Results

4.1 Growth and temporal development of the field

The corpus demonstrates an unusually strong recent acceleration. Publication activity remained extremely limited during the first decade: only one document was recorded in each of 2004 and 2005, followed by intermittent output through 2017. Annual production increased to nine publications in 2018, 13 in both 2019 and 2020, 16 in 2021, 28 in 2022, and 30 in 2023. A structural break becomes visible after 2023. Output increased to 70 publications in 2024 and then to 179 in 2025. The supplied 2026 corpus already contains 170 publications, despite 2026 being an incomplete year. Of the 545 documents, 528 (96.9%) were published from 2018 onward, while 419 (76.9%) appeared during 2024–2026 alone. Furthermore, 349 documents (64.0%) belong to 2025–2026. These figures demonstrate that AI in real-time payments is not a mature field characterized by steady long-term accumulation. Rather, it is a rapidly expanding intersection whose contemporary literature has formed predominantly during the past three years. The reported annual growth rate of 26.29% reinforces this conclusion.




Figure 1. Annual scientific production on artificial intelligence in real-time payment systems, 2004–2026.

Source: Authors' analysis based on Scopus bibliographic data using Bibliometrix/Biblioshiny in R (n = 545 documents). Note: The 2026 value represents a partial publication year and should not be interpreted as a full-year total.

4.2 Publication outlets and disciplinary dispersion

The most relevant source is Lecture Notes in Networks and Systems, with 21 publications, followed by the expanded Lecture Notes in Computer Science series with 16 publications. IEEE Access ranks third with ten publications, while the IEEE Internet of Things Journal and Procedia Computer Science contribute seven documents each. Other active outlets include Communications in Computer and Information Science, ACM conference proceedings, Journal of Risk and Financial Management, Expert Systems with Applications, IEEE Transactions on Information Forensics and Security, and Journal of Digital Banking. The distribution is highly fragmented. The ten most productive sources collectively contribute only approximately 15.4% of the corpus, while 402 sources collectively host 545 documents. This absence of a dominant outlet suggests that the field has not yet consolidated around a small group of specialist journals. The finding is consistent with the interdisciplinary character of the topic. Technical studies appear in AI, computer science, IoT, cybersecurity, and information-systems outlets, whereas organizational and financial studies are distributed across banking, FinTech, risk, management, and finance journals. The dominance of conference papers—296 documents or 54.3%—is equally significant. In fast-developing computing disciplines, conference publication can accelerate dissemination of new

algorithms and architectures. Within this corpus, the conference-heavy profile suggests that the technical frontier is advancing more rapidly than journal-based theoretical consolidation.




Figure 2. Most relevant publication sources in artificial intelligence and real-time payment systems research.

Source: Authors' analysis based on Scopus bibliographic records using Bibliometrix/Biblioshiny in R ($n = 545$ documents; 2004–2026). Publication sources are ranked according to the number of documents in the analyzed corpus.

4.3 Authors and institutional structure

No author dominates the field. A. Kumar and X. Li each contribute six publications, while S. Gupta, S. Kumar, K. Sandkuhl, Y. Wang, and Z. Zhang each contribute five. Several authors contribute four or three publications. Fractionalized authorship provides additional insight. Sandkuhl's five documents yield a fractionalized contribution of 2.67, substantially higher than several other authors with five or six papers. This indicates a comparatively larger individual contribution within smaller author teams. His prominence is also substantively consistent with early work explicitly connecting fraud detection with instant-payment requirements (Diadiushkin et al., 2019). The institutional structure is similarly dispersed. SRM Institute of Science and Technology records seven publications, followed by IEC College of Engineering and Technology, the University of Patras, and the University of the Cumberlands, with six each. Several institutions—including Beijing University of Posts and Telecommunications, King Saud University, Qatar University, Southeast University, Tsinghua University, and the University of Chinese Academy of Sciences—record five publications. Eight records are classified as NOTREPORTED, demonstrating an affiliation-metadata limitation. The absence of a single institution with substantially greater output than the others suggests that the field is developing through distributed project-level research rather than being concentrated within one established intellectual center.

4.4 Geographical distribution and collaboration

India has the highest country-affiliation frequency (437), followed by the United States (210) and China (177). The United Kingdom records 39 affiliations, followed by Saudi Arabia (32), Iraq (27), Uzbekistan (25), South Korea (23), Morocco (19), and Pakistan (17). Bangladesh and the United Arab Emirates each record 16. These numbers should not be interpreted as mutually exclusive document counts, because international publications can contribute affiliations to several countries. Nevertheless, they establish a clear geographical pattern: research

production is strongly concentrated in India, the United States, and China, with substantial participation from emerging Asian and Middle Eastern research systems.

Figure 3. Geographical structure of research on artificial intelligence in real-time payment systems





Figure 3(a): country scientific production



3(b) international research collaboration network.

Source: Authors' analysis based on Scopus bibliographic data using Bibliometrix/Biblioshiny in R ($n = 545$ documents; 2004–2026). Note: Country productivity is based on author-affiliation information; internationally co-authored publications may therefore contribute to more than one country. Collaboration links represent cross-country co-authorship relationships.

The collaboration network shows that the India–USA relationship is the strongest bilateral link, with 22 collaborative occurrences. India–Saudi Arabia and USA–China each record eight, followed by India–Uzbekistan and USA–Bangladesh with five each. China–Hong Kong, China–United Kingdom, India–Canada, India–Oman, India–United Arab Emirates, Saudi Arabia–Malaysia, and United Arab Emirates–Malaysia record four each. Although international collaboration is visible, the overall international co-authorship rate is only 25.32%. Thus approximately three quarters of documents are not classified as internationally co-authored. For a problem involving globally connected payment infrastructure, cybercrime networks, interoperable financial standards, and cross-border fraud, this level of international collaboration leaves substantial room for deeper comparative and cross-jurisdictional research.

4.5 Citation structure

The corpus has an average of 9.415 citations per document, although citation distribution is strongly skewed. The highest globally cited record is a 2019 IEEE Transactions on Industrial Informatics publication with 385 citations. Roy et al.'s (2018) deep-learning fraud-detection study records 268 citations, while Allen et al.'s (2021) FinTech survey records 252. Other highly cited records concern IoT, blockchain marketplaces, distributed systems,

4.6 Keyword structure: the conceptual core of the literature

The terms reveal three interacting conceptual layers. The first is the detection layer, represented by fraud detection, anomaly detection, crime, random forests, decision trees, neural networks, XGBoost, supervised and unsupervised learning, and deep learning. The second is the payment and financial layer, represented by electronic money, financial transactions, real-time transactions, payment systems, banking, unified payment interface (UPI), and digital payments. The third is the infrastructure and governance layer, represented by blockchain, network security, cybersecurity, data privacy, authentication, regulatory compliance, distributed systems, smart contracts, and decentralized finance. The emergence of explainable AI, adversarial machine learning, real-time transaction monitoring, and graph neural networks additionally indicates that the research agenda is expanding beyond basic predictive classification toward operational robustness and accountability.

6246

Source: Authors' analysis based on Scopus bibliographic data using Bibliometrix/Biblioshiny in R (KW_Merged; 250 terms; minimum frequency = 2; Louvain clustering). Node communities represent conceptually related keyword groups, while links indicate co-occurrence relationships.

The co-word network contains 49 nodes distributed across three communities. Centrality metrics identify “learning systems,” “crime,” and “fraud detection” as among the most influential terms in the primary cluster. Machine learning, deep learning, anomaly detection, financial transactions, and electronic money form closely connected concepts. The second major cluster centers on artificial intelligence, blockchain, and network security, with blockchain showing particularly strong bridging behavior. This cluster links AI to cybersecurity, privacy, distributed infrastructures, and financial-system architecture. A much smaller third community contains reinforcement learning terminology. Although its current network size is limited, its separation from the dominant fraud-detection cluster suggests a specialist trajectory that has not yet fully integrated with mainstream RTP research. Conceptually, the co-word network therefore shows that the field is organized around two large poles: Firstly, AI as a transaction-risk and fraud-detection capability, and second, AI as part of a secure, distributed, digitally mediated payment infrastructure. The RTP domain lies at the intersection of these poles.

4.8 Temporal evolution of research topics




Figure 5. Evolution of dominant and emerging research topics in artificial intelligence-enabled real-time payment systems.

Source: Authors' analysis based on Scopus bibliographic data using Bibliometrix/Biblioshiny in R (n = 545 documents; 2004–2026). Bubble size reflects the relative frequency of each topic, while its temporal position indicates the period in which the topic was most prominent.

Trend-topic analysis provides particularly strong evidence of thematic transformation. The earliest clearly RTP-specific trending term is “instant payment,” whose first quartile year is 2019, median year 2020, and third quartile year 2022. This suggests that explicit concern with instant-payment environments became visible before the recent surge in AI literature. A second stage appears between 2021 and 2025. Internet of Things, distributed ledger, scalability, bitcoin, blockchain, and security become prominent. Blockchain has a frequency of 146 and a median trending year of 2024, indicating that distributed financial architectures became one of the main conceptual

bridges into the present literature. The most recent stage is markedly different. Learning systems (176), crime (175), and fraud detection (152) all have median years of 2025 and third-quartile years extending into 2026. Graph neural networks (51) have a median year of 2026, as do behavioral research and financial-fraud-detection variants. The literature is no longer primarily asking whether digital payments can be made immediate. It is increasingly asking how immediate payments can remain secure when risk decisions must be made in milliseconds and adversaries adapt dynamically.

4.9 Thematic map

The thematic map identifies four major clusters.




Figure 6. Strategic thematic map of artificial intelligence research in real-time payment systems based on Callon centrality and density.

Source: Authors' analysis based on Scopus bibliographic data using Bibliometrix/Biblioshiny in R (KW_Merged; 250 terms; minimum frequency = 2; Louvain clustering). The horizontal axis represents Callon centrality and the vertical axis represents Callon density.

Note: The upper-right quadrant represents motor themes; the upper-left quadrant represents highly developed but relatively isolated niche themes; the lower-right quadrant represents basic or transversal themes; and the lower-left quadrant contains weakly developed themes that may be emerging or declining.

Table 2. Strategic thematic structure

Theme	Callon centrality	Callon density	Cluster frequency	Interpretation
Learning systems / fraud detection	11.06	28.72	2,156	Motor theme
Artificial intelligence / blockchain / security	9.60	27.17	1,426	Motor theme
Behavioral research / FinTech analytics	6.99	25.51	589	Emerging/developing

Graph neural networks / reinforcement learning	3.51	22.04	280	Emerging frontier
--	------	-------	-----	-------------------

Source: Authors' analysis of the supplied Biblioshiny thematic map.

4.9.1 Learning systems and fraud detection

The largest and most central cluster combines learning systems, crime, fraud detection, machine learning, deep learning, finance, anomaly detection, electronic money, real-time transactions, random forests, adversarial machine learning, logistic regression, decision trees, explainable AI, real-time fraud detection, transaction monitoring, UPI, and digital payments. With a Callon centrality of 11.06, density of 28.72, and cluster frequency of 2,156, it is both the most connected and internally developed theme. It can therefore be interpreted as the field's strongest motor theme. Importantly, the cluster combines established algorithms with newer concerns. Conventional logistic regression, decision trees, random forests, and supervised learning coexist with deep neural networks, adversarial learning, explainable AI, and real-time monitoring. The cluster thus captures the historical transformation of fraud analytics rather than one specific algorithmic paradigm.

4.9.2 Artificial intelligence, blockchain, and security

The second motor theme has a centrality of 9.60, density of 27.17, and frequency of 1,426. Its major terms include artificial intelligence, blockchain, network security, data privacy, cybersecurity, scalability, decentralized finance, distributed systems, regulatory compliance, authentication, and smart contracts. The position of this cluster demonstrates that payment intelligence cannot be analytically separated from payment infrastructure. Research increasingly treats fraud detection, cybersecurity, privacy, distributed trust, and regulatory compliance as components of the same system.

4.9.3 Graph neural networks and reinforcement learning

The graph-neural-network cluster has lower centrality (3.51) and density (22.04) but includes graph neural networks, reinforcement learning, deep reinforcement learning, heterogeneous graphs, benchmarking, phishing, fraud, and instant payment. A low-centrality/low-density theme could theoretically represent either an emerging or declining topic. Temporal evidence resolves this ambiguity: graph neural networks have a median trending year of 2026, and recent highly coupled papers explicitly address real-time graph-based transaction fraud. The cluster is therefore best interpreted as an emerging frontier, not a declining one.

4.9.4 Behavioral research and FinTech analytics

The fourth cluster has centrality 6.99, density 25.51, and frequency 589. It includes behavioral research, banking, big data, decision-making, FinTech, intelligent systems, e-commerce, information management, predictive analytics, risk management, investment, digital transformation, and financial markets. This cluster broadens the field beyond engineering. Its emergence suggests growing recognition that AI-enabled payment security is embedded within organizational decision-making, customer behavior, financial institutions, governance, and risk management. The cluster remains less central than the technical motor themes, suggesting that social, organizational, and behavioral research has not yet fully integrated with the engineering core.

4.10 Bibliographic coupling and contemporary research fronts

Document-level bibliographic coupling identifies three distinct research fronts.

Table 3. Bibliographic-coupling research fronts

Cluster	Main descriptors	Documents	Centrality	Local impact
1	Fraud detection, machine learning, anomaly detection	91	0.356	6.109
2	Machine learning, fraud detection, deep learning	58	0.450	1.753
3	Blockchain, security, machine learning	101	0.411	1.121

Source: Authors' analysis of the supplied Biblioshiny coupling map.

Cluster 1 has the highest local impact. Particularly influential recent documents include Li and Chen's (2025) dynamic real-time banking anti-fraud model, Nguyen and Le's (2025) heterogeneous temporal GNN study, and Lu et al.'s (2022) BRIGHT framework. The concentration of these studies confirms that operational real-time fraud detection has become a major contemporary research front. Cluster 2 achieves the highest coupling centrality (**0.450**), suggesting that machine-learning and deep-learning fraud detection draws on a broad shared knowledge base linking several areas of the field. Cluster 3 is the largest, with **101 documents**, and is dominated by blockchain and security. Its lower local impact relative to Cluster 1 suggests that blockchain/security represents a broad architectural stream, whereas the more narrowly defined real-time fraud-detection stream currently contains particularly influential locally connected publications.

5. Discussion

5.1 From payment speed to intelligent risk control

The central finding of this study is not merely that publications have increased. The more important result is a change in the object of research. Earlier work focused on the technological possibility and operational implications of instant payment. The literature then expanded into blockchain, distributed systems, IoT, scalability, and security. The newest literature places AI-enabled fraud detection at the center and increasingly asks how learning systems can operate under real-time constraints. This progression corresponds closely to the operational evolution of fast-payment infrastructures. Once immediate payments become technically feasible and broadly accessible, the limiting problem shifts from transaction speed to safe transaction speed. The result also explains why fraud detection has become the dominant motor theme. In real-time payment environments, the fraud model is no longer a peripheral monitoring tool operating after transaction processing; it increasingly becomes part of the transaction-decision infrastructure itself.

5.2 Accuracy alone is becoming an inadequate research objective

Many earlier financial-fraud studies were structured as model-comparison exercises in which competing classifiers were evaluated primarily using predictive performance. Such studies remain valuable, but RTP deployment imposes requirements that conventional benchmark accuracy does not capture. A model with marginally higher F1 score may be operationally inferior if it requires hundreds of milliseconds of additional inference time, depends on features unavailable at authorization, produces unstable predictions under concept drift, or generates explanations that compliance teams cannot use.

The bibliometric emergence of terms including “real-time transaction monitoring,” “explainable AI,” “adversarial machine learning,” “data privacy,” “regulatory compliance,” and “benchmarking” reflects this broadening of evaluation criteria. This shift is visible in recent real-time fraud studies. Li and Chen (2025) explicitly develop an anti-fraud model around real-time transaction flows, while BRIGHT is architected around online inference efficiency rather than offline predictive quality alone (Lu et al., 2022).

5.3 Graph neural networks represent the strongest emerging technical frontier

The combination of trend analysis, thematic mapping, and bibliographic coupling provides unusually consistent evidence regarding graph learning. First, “graph neural networks” appears 51 times and has a median trend year of 2026. Second, GNNs form a separate thematic cluster alongside reinforcement learning, benchmarking, heterogeneous graphs, and instant payment. Third, two of the strongest recent coupled documents—Lu et al. (2022) and Nguyen and Le (2025)—specifically use graph architectures for real-time fraud detection. This convergence suggests that GNNs are not simply another classifier entering an already crowded fraud literature. They respond to a structural limitation of transaction-level models: fraudulent behavior is relational. Nevertheless, graph methods introduce unresolved problems. Financial graphs can contain millions of rapidly changing nodes and edges. Neighborhood aggregation can increase inference cost. Temporal leakage can produce unrealistic evaluation results. Fraudulent accounts may deliberately alter graph structure to avoid detection. Labels remain imbalanced, and graph explanations can be difficult for compliance investigators to interpret. Motie and Raahemi's (2024) systematic review reinforces the importance of these challenges across financial-fraud GNN research. The

next phase should therefore move beyond demonstrating that GNNs outperform tabular baselines and address whether they can be safely and efficiently deployed in regulated real-time payment infrastructures.

5.4 Blockchain is important, but its role requires conceptual clarification

Blockchain is among the most frequent concepts in the corpus and anchors the second motor theme. This prominence demonstrates the strong intellectual connection between AI, distributed financial systems, and security. However, bibliometric frequency should not be interpreted as evidence that blockchain is necessarily a core component of all real-time payment systems. Many national fast-payment infrastructures operate through centralized or hybrid architectures rather than public blockchains. The result should therefore be interpreted at a higher conceptual level. Blockchain research contributes to the AI–RTP knowledge base through problems of distributed trust, authentication, smart contracts, privacy, decentralized finance, tamper resistance, and secure digital value transfer. These topics overlap with RTP security even when the underlying infrastructures differ. Future reviews should distinguish more carefully among real-time retail payment systems, blockchain payments, cryptocurrencies, decentralized finance, and real-time gross settlement. Treating all of these as interchangeable could obscure rather than clarify the field.

5.5 The behavioral and governance dimension remains comparatively underdeveloped

The behavioral-research cluster is recent but less central than the two technical motor themes. This asymmetry is consequential. An AI fraud-detection system does not act in isolation. False positives affect customers; automated account blocking can create exclusion risks; explainability affects investigators and regulators; customer behavior changes in response to friction; and attackers adapt to institutional controls. Organizations must decide thresholds, escalation procedures, human-review policies, liability arrangements, and acceptable trade-offs between convenience and security. Yet the bibliometric structure suggests that technical research has advanced more quickly than research integrating these organizational and behavioral questions. This is consistent with the gap identified by broader bibliometric work on AI and financial fraud, which points toward greater need for ethical, organizational, and interdisciplinary consideration (Moura et al., 2025). A mature RTP research field will therefore require stronger integration between computer science, financial regulation, behavioral science, risk management, information systems, and payment economics.

5.6 Geography and collaboration reveal an opportunity for comparative research

India, the United States, and China dominate country-affiliation frequencies. India is particularly prominent, and India–USA represents the strongest international collaboration dyad. The geographical pattern is substantively meaningful because these countries represent large digital-payment ecosystems and substantial AI research capacity. At the same time, the field includes visible contributions from Saudi Arabia, the United Arab Emirates, Pakistan, Bangladesh, Iraq, Morocco, Uzbekistan, Malaysia, and other emerging markets. This distribution creates an opportunity that current collaboration patterns have not fully exploited. Fraud characteristics, digital identities, payment architectures, consumer behaviors, and regulatory rules differ considerably across jurisdictions. Yet only 25.32% of publications involve international co-authorship. Cross-country research would allow scholars to determine whether AI fraud models generalize across payment systems or merely perform well within the institutional environment from which their data were collected.

6. Future Research Agenda

The synthesis supports seven priority areas.

Table 4. Proposed future research agenda for AI in real-time payment systems

Priority	Core research question	Recommended direction
Temporal and graph intelligence	How can relational fraud patterns be identified without temporal leakage?	Temporal GNNs, heterogeneous graphs, causal-time splits

Concept drift	How quickly do fraud models become obsolete?	Online learning, drift detection, continual learning
Operational benchmarking	Can models meet production latency and throughput requirements?	P95/P99 latency, throughput, memory and computational-cost reporting
Explainability	Can payment-risk decisions be understood and contested?	Local explanations, graph explanations, investigator-centered XAI
Privacy and collaboration	Can institutions learn jointly without centralizing sensitive data?	Federated learning, privacy-preserving analytics, secure aggregation
Adversarial resilience	How robust are models to adaptive fraudsters?	Adversarial ML, graph manipulation tests, red-team evaluation
Governance and societal impact	How should automated RTP intervention be governed?	Fairness, accountability, human oversight, cross-jurisdictional regulation

Source: Authors' analysis

6.1 Temporal and heterogeneous graph learning

Future research should emphasize dynamic rather than static transaction graphs. RTP networks change continuously, and models must distinguish relationships available at prediction time from information generated afterward. Evaluation should therefore use time-respecting data splits and explicitly test for future-information leakage. Heterogeneous networks containing accounts, devices, merchants, beneficiaries, locations, and payment instruments may be particularly useful because fraudulent activity frequently spans entity types. Recent real-time heterogeneous GNN research indicates that this trajectory is already emerging (Nguyen & Le, 2025).

6.2 Concept drift and continual learning

The dominant literature remains heavily focused on training and testing models on fixed historical datasets. This is insufficient for adversarial environments in which fraud strategies change in response to detection. Future studies should report performance degradation over time, drift-detection procedures, retraining frequency, label-delay sensitivity, and catastrophic-forgetting risks. Research should also distinguish gradual behavioral drift from abrupt changes associated with new fraud campaigns.

6.3 Explainable AI and human-in-the-loop investigation

Explainable AI appears in the motor fraud-detection cluster but remains less frequent than core predictive terms. This imbalance should narrow. Payment institutions must often explain why a transaction was delayed, rejected, escalated, or associated with suspicious behavior. Future XAI research should therefore evaluate explanations not simply by technical fidelity but by whether fraud analysts, compliance personnel, auditors, regulators, and customers can meaningfully use them. Graph systems create a particular challenge because decisions may depend on multi-hop relationships rather than easily interpretable transaction attributes.

6.4 Privacy-preserving and federated intelligence

Fraud patterns often span financial institutions, but privacy, confidentiality, competition, and regulation constrain direct sharing of customer-level transaction data. Federated learning and related privacy-preserving techniques could allow institutions to learn from distributed fraud signals without pooling raw records. However, this research should move beyond proof-of-concept accuracy. Studies must examine communication overhead, poisoning

attacks, non-identically distributed institutional data, privacy leakage, governance of shared models, and responsibility when collaborative models produce harmful decisions.

6.5 Adversarial machine learning

Fraud detection differs from many conventional predictive tasks because the predicted subjects actively attempt to defeat the classifier. The presence of adversarial machine learning within the principal motor theme suggests growing recognition of this fact. Future work should evaluate not only clean-data performance but also robustness to manipulated transaction attributes, synthetic identities, account networks, graph-edge manipulation, coordinated mule activity, and model-probing behavior. Security evaluation should therefore become an integral component of AI-payment benchmarking.

6.6 Responsible AI, regulation, and institutional governance

Technical capability alone will not determine successful AI adoption in RTP. Researchers should examine who is accountable when an algorithm wrongly blocks legitimate payments, how customers can contest automated decisions, what level of explanation regulators require, whether model performance differs systematically among user groups, how long fraud-related data should be retained, and when human intervention should override automated decisions. The convergence between AI, data privacy, regulatory compliance, authentication, and behavioral research identified in the thematic map suggests that these concerns will become increasingly central as the field matures.

7. Theoretical, Managerial, and Policy Implications

7.1 Theoretical implications

The analysis suggests that AI in real-time payments should be conceptualized as a socio-technical risk-control system, not merely a machine-learning application. The literature presently contains three partially overlapping levels:

1. Transaction intelligence, concerned with detecting anomalous or fraudulent behavior;
2. Infrastructure intelligence, concerned with security, blockchain, authentication, and distributed systems; and
3. Institutional intelligence, concerned with banking decisions, risk management, regulation, behavior, and governance.

Future theory should integrate these levels. A technically accurate model can still fail if it is too slow, operationally expensive, legally opaque, behaviorally disruptive, or incompatible with the architecture of the payment system. This provides a broader conceptual foundation for research: effective RTP intelligence is produced by the interaction of predictive capability, temporal responsiveness, network awareness, institutional governance, and human decision-making.

7.2 Implications for financial institutions and FinTech firms

For banks, payment-service providers, and FinTech firms, the results indicate that investing solely in increasingly complex predictive algorithms may deliver diminishing returns if underlying data pipelines and deployment architecture remain weak. Operational systems should combine streaming feature engineering, model monitoring, graph intelligence, rule-based safeguards, explainability, human escalation, and continual updating. Institutions should also evaluate models using monetary loss and customer-friction metrics rather than conventional accuracy alone. A small improvement in recall may have substantial financial value when it identifies high-value fraud, whereas a marginally more accurate model may be commercially harmful if false positives inconvenience large numbers of legitimate customers.

7.3 Implications for regulators and payment-system operators

Regulators and payment operators face a difficult balance. Stronger AI controls may reduce fraud but can increase opacity and automate consequential financial decisions. Regulatory frameworks should therefore encourage

measurable model governance, including documentation of training data, performance monitoring, drift management, explanation procedures, audit trails, human escalation, cybersecurity testing, and accountability for automated interventions. Cross-institution collaboration is particularly important because payment fraud frequently moves across institutional boundaries. Regulators and payment operators may be able to facilitate privacy-preserving information-sharing frameworks without requiring unrestricted exchange of customer-level data.

8. Limitations

First, bibliometric findings are bounded by the search-defined corpus. Bibliometric analysis cannot recover relevant studies excluded by the original query and may include peripheral publications captured through broad terminology. Second, the supplied Biblioshiny report does not retain the exact Scopus search syntax, extraction date, or detailed screening log. These should be restored from the original Scopus search before journal submission to ensure full reproducibility. Third, the corpus displays evidence of semantic spillover. Some highly cited records and trend terms arise from adjacent IoT, energy, blockchain, and general AI literatures rather than narrowly defined RTP applications. This interdisciplinary overlap is substantively informative but also indicates that manual title/abstract relevance screening would strengthen corpus precision. Fourth, keyword normalization is incomplete. Variants such as “machine learning” and “machine-learning,” “blockchain” and “block-chain,” and different representations of “real time” remain separate. A controlled thesaurus should therefore be applied before producing final publication figures. Fifth, citation analysis is subject to citation-age bias. The extremely recent 2025–2026 literature has had limited time to accumulate citations. Sixth, 2026 represents a partial publication year, making its raw document count inappropriate for direct full-year comparison with 2025. Seventh, bibliographic metadata capture scientific attention but not necessarily technological adoption. A highly central research theme may not yet be widely deployed in real-world payment infrastructure. Finally, bibliometric mapping reveals structural associations rather than causal relationships. Keyword co-occurrence, collaboration, and bibliographic coupling indicate intellectual proximity, not proof that one technology causes another research trajectory.

9. Conclusion

This study maps the emerging research domain at the intersection of artificial intelligence and real-time payment systems using a Bibliometrix/Biblioshiny analysis of 545 documents published between 2004 and 2026. The field is characterized by exceptionally rapid recent expansion. Its 26.29% annual growth rate, concentration of 76.9% of publications in 2024–2026, and dominance of conference publications demonstrate a technically dynamic domain whose scholarly structure is still consolidating. More importantly, the bibliometric evidence reveals a clear thematic evolution. Early interest in instant payment is followed by research on distributed ledgers, IoT, blockchain, scalability, and security. The most recent literature shifts strongly toward learning systems, fraud detection, machine learning, deep learning, and graph neural networks. Two established motor themes now structure the field: AI-enabled fraud analytics and AI–blockchain–security infrastructure. Meanwhile, graph neural networks and reinforcement learning form a distinctly recent emerging frontier, and behavioral research, FinTech analytics, risk management, and institutional decision-making signal an expanding socio-technical research agenda. Bibliographic coupling reinforces this transformation. Operational fraud and anomaly detection exhibit the highest local impact; machine/deep-learning research occupies the most structurally central coupled position; and blockchain/security forms the largest architectural stream. Recent influential work on real-time ML and temporal GNNs indicates that the research frontier is moving from static transaction classification toward relational and temporally aware intelligence. The next stage of scholarship should therefore move beyond the question of which algorithm detects fraud most accurately. The more consequential question is: How can an AI system make secure, explainable, privacy-preserving, robust, and economically defensible decisions within the milliseconds available in a real-time payment environment? Answering that question requires integration of temporal graph learning, streaming analytics, concept-drift adaptation, adversarial robustness, federated intelligence, operational latency measurement, explainability, human oversight, and responsible AI governance. The long-term development of real-time payment intelligence will depend not on AI alone, but on the ability to

align algorithmic intelligence with payment infrastructure, institutional responsibility, and the continuously changing behavior of legitimate users and adversaries.

Conflict of Interest Statement

The author(s) declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

Ethical Approval Statement

Ethical approval was not required for this study because it is based exclusively on bibliographic metadata retrieved from the Scopus database and does not involve human participants, animals, clinical interventions, or personally identifiable information.

Data Availability Statement

The bibliographic data analyzed in this study were retrieved from the Scopus database. The bibliometric analyses were conducted using the Bibliometrix package and Biblioshiny interface in R. The underlying bibliographic dataset is subject to the access and licensing conditions of Scopus. Data supporting the bibliometric results may be made available by the corresponding author upon reasonable request, subject to applicable database licensing restrictions.

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] Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40, 100402. <https://doi.org/10.1016/j.cosrev.2021.100402>
- [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] Baesens, B., Höppner, S., & Verdonck, T. (2021). Data engineering for fraud detection. *Decision Support Systems*, 150, 113492. <https://doi.org/10.1016/j.dss.2021.113492>
- [6] Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182–194. <https://doi.org/10.1016/j.inffus.2017.09.005>
- [7] Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331. <https://doi.org/10.1016/j.ins.2019.05.042>
- [8] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357. <https://doi.org/10.1613/jair.953>
- [9] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939785>
- [10] Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University—Computer and Information Sciences*, 35, 145–174. <https://doi.org/10.1016/j.jksuci.2022.11.008>
- [11] Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University—Computer and Information Sciences*, 35, 145–174. <https://doi.org/10.1016/j.jksuci.2022.11.008>

- [12] Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). Science mapping software tools: Review, analysis, and cooperative study among tools. *Journal of the American Society for Information Science and Technology*, 62(7), 1382–1402. <https://doi.org/10.1002/asi.21525>
- [13] Committee on Payments and Market Infrastructures. (2016). Fast payments—Enhancing the speed and availability of retail payments (CPMI Report No. 154). Bank for International Settlements.
- [14] Committee on Payments and Market Infrastructures. (2016). Fast payments—Enhancing the speed and availability of retail payments (CPMI Report No. 154). Bank for International Settlements.
- [15] Committee on Payments and Market Infrastructures. (2021). Developments in retail fast payments and implications for RTGS systems (CPMI Papers No. 201). Bank for International Settlements.
- [16] 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). Bank for International Settlements.
- [17] Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784–3797. <https://doi.org/10.1109/TNNLS.2017.2736643>
- [18] Diadiushkin, A., Sandkuhl, K., & Maiatin, A. (2019). Fraud detection in payments transactions: Overview of existing approaches and usage for instant payments. *Complex Systems Informatics and Modeling Quarterly*, 20, 72–88. <https://doi.org/10.7250/csimq.2019-20.04>
- [19] Diadiushkin, A., Sandkuhl, K., & Maiatin, A. (2019). Fraud detection in payments transactions: Overview of existing approaches and usage for instant payments. *Complex Systems Informatics and Modeling Quarterly*, 20, 72–88. <https://doi.org/10.7250/csimq.2019-20.04>
- [20] 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>
- [21] Ellegaard, O., & Wallin, J. A. (2015). The bibliometric analysis of scholarly production: How great is the impact? *Scientometrics*, 105(3), 1809–1831. <https://doi.org/10.1007/s11192-015-1645-z>
- [22] Frost, J., Koo Wilkens, P., Kosse, A., Shreeti, V., & Velásquez, C. (2024). Fast payments: Design and adoption. *BIS Quarterly Review*, March 2024. Bank for International Settlements.
- [23] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), Article 44. <https://doi.org/10.1145/2523813>
- [24] Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429. <https://doi.org/10.1016/j.eswa.2021.116429>
- [25] Höppner, S., Baesens, B., Verbeke, W., & Verdonck, T. (2022). Instance-dependent cost-sensitive learning for detecting transfer fraud. *European Journal of Operational Research*, 297(1), 291–300. <https://doi.org/10.1016/j.ejor.2021.05.028>
- [26] Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234–245. <https://doi.org/10.1016/j.eswa.2018.01.037>
- [27] Li, F., & Chen, Z. (2025). Dynamic quantification anti-fraud machine learning model for real-time transaction fraud detection in banking. *Discover Computing*, 28, 59. <https://doi.org/10.1007/s10791-025-09549-7>
- [28] Lim, W. M., Kumar, S., & Donthu, N. (2024). How to combine and clean bibliometric data and use bibliometric tools synergistically: Guidelines using metaverse research. *Journal of Business Research*, 182, 114760. <https://doi.org/10.1016/j.jbusres.2024.114760>
- [29] Lu, M., Han, Z., Rao, S. X., Zhang, Z., Zhao, Y., Shan, Y., Raghunathan, R., Zhang, C., & Jiang, J. (2022). BRIGHT—Graph neural networks in real-time fraud detection. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (pp. 3342–3351). Association for Computing Machinery. <https://doi.org/10.1145/3511808.3557136>
- [30] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 4765–4774). Curran Associates.

- [31] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (Vol. 54, pp. 1273–1282). *Proceedings of Machine Learning Research*.
- [32] 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>
- [33] Moura, L., Barcaui, A., & Payer, R. (2025). AI and financial fraud prevention: Mapping the trends and challenges through a bibliometric lens. *Journal of Risk and Financial Management*, 18(6), 323. <https://doi.org/10.3390/jrfm18060323>
- [34] Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569. <https://doi.org/10.1016/j.dss.2010.08.006>
- [35] Nguyen, H., & Le, B. (2025). Real-time transaction fraud detection via heterogeneous temporal graph neural network. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025)* (Vol. 2, pp. 364–375). SCITEPRESS. <https://doi.org/10.5220/0013161300003890>
- [36] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939778>
- [37] 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>
- [38] Tanwar, S., & Arora, R. G. (2025). Two decades of transformation: A bibliometric journey through the evolution and trends in digital payment systems. *Global Knowledge, Memory and Communication*. Advance online publication. <https://doi.org/10.1108/GKMC-08-2024-0531>
- [39] Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing*, 149, 110984. <https://doi.org/10.1016/j.asoc.2023.110984>
- [40] Van Belle, R., Baesens, B., & De Weerd, J. (2023). CATCHM: A novel network-based credit card fraud detection method using node representation learning. *Decision Support Systems*, 164, 113866. <https://doi.org/10.1016/j.dss.2022.113866>
- [41] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66. <https://doi.org/10.1016/j.cose.2015.09.005>
- [42] Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429–472. <https://doi.org/10.1177/1094428114562629>
- [43] Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429–472. <https://doi.org/10.1177/1094428114562629>