

Artificial Intelligence in Auditing and Banking Risk Management: A Bibliometric Mapping of Explainability, Fraud Detection, and Regulatory Compliance

Hicham BAHIDA, (Lecturer-Researcher in Economics and Management)

National School of Business and Management of Dakhla (ENCG Dakhla)

Ibn Zohr University, Dakhla, Morocco

Fatima AARAB, (PhD in Economics and Management)

Faculty of Economics and Management

Ibn Zohr University, Agadir, Morocco

Correspondence address:	Fatima AARAB, Faculty of Economics and Management, Ibn Zohr University, Agadir, Morocco
Disclosure Statement:	The authors declare that they have not received any financial support that could have influenced the objectivity of this study. They take full responsibility for any potential plagiarism, the use of artificial intelligence in the writing process, as well as for the results presented in this article.
Conflict of Interest:	The authors report no conflicts of interest.
Cite this article:	BAHIDA, H., & AARAB, F. (2026). Artificial Intelligence in Auditing and Banking Risk Management: A Bibliometric Mapping of Explainability, Fraud Detection, and Regulatory Compliance. <i>International Journal of Accounting, Finance, Auditing, Management and Economics</i> , 7(13), 362–383. https://doi.org/10.5281/zenodo.23241806
License	This is an open access article under the CC BY-NC-ND license

Received: 12/09/2026

Accepted: 16/10/2026

International Journal of Accounting, Finance, Auditing, Management and Economics - IJAFAME

ISSN: 2658-8455

Volume 7, Issue 13 (2026)

Artificial Intelligence in Auditing and Banking Risk Management: A Bibliometric Mapping of Explainability, Fraud Detection, and Regulatory Compliance

Abstract:

Artificial intelligence is already present in a wide range of banking and audit activities, but the literature around these uses is still split across several research conversations. Banking studies usually concentrate on prediction, fraud monitoring, cybersecurity, and risk management. Audit research asks a different question: can the outputs be documented, challenged, controlled, and used in professional judgment? Work on explainable AI adds another layer, namely whether complex model outputs can be understood well enough to support scrutiny. We focus on the overlap among these streams rather than on AI in banking as a whole. To date, this overlap has not been the subject of a dedicated bibliometric mapping, which is the gap this study addresses. Four questions guide the analysis: how have scientific production and influence evolved, which concepts structure the field, which references does it draw on, and how do authors and countries collaborate? The dataset contains 43 Scopus-indexed documents published between 2014 and 2026. Using Bibliometrix/Biblioshiny in R, we combined performance analysis, thematic mapping and evolution, multiple correspondence analysis (MCA), co-citation, historiography, and collaboration networks. Screening was intentionally selective. Prediction-only papers were excluded when assurance, control, auditability, or related oversight issues were not substantive. Thirty-four papers (79.1%) were published between 2024 and 2026. AI and machine learning form the technical core of the literature, with explainability, financial risk management, fraud detection, cybersecurity, and regulatory compliance recurring around it. MCA identifies three clusters: a broad AI-risk-fraud-explainability cluster, a smaller audit-compliance-fintech cluster, and a specialized blockchain-regulatory-compliance cluster. Citation and collaboration patterns are less settled: co-citation communities remain fragmented and direct citation chains are short. The literature therefore looks less like a finished theoretical framework than a set of technical and governance concerns that are beginning to converge. We interpret this convergence through the notion of AI-enabled financial risk governance and assurance, which distinguishes AI capability from the assurance capability needed to audit and control its use. In that setting, explainability and auditability matter because even a high-performing model still has to withstand professional challenge and regulatory scrutiny.

Keywords: Artificial Intelligence; Auditing; Banking; Bibliometric Analysis; Explainable Artificial Intelligence (XAI); Financial Risk Management; Regulatory Compliance.

JEL Classification: G21; G32; M42; O33

Paper type: Empirical Research (bibliometric study)

1. Introduction

Digital banking predates the current surge of interest in AI. Electronic channels and mobile services first changed how customers reached financial institutions; later, data-intensive systems began to reshape how those institutions assessed customers and transactions. AI and machine learning extend this development because their role goes beyond routine automation. They classify cases, estimate risk, detect unusual patterns, and support decisions. In financial services, these applications now include credit assessment, fraud detection, anti-money-laundering controls, customer analytics, cybersecurity, and risk management (Königstorfer & Thalmann, 2020; Pattnaik et al., 2024). This agenda grew out of an earlier literature on technology adoption, in which the use of mobile banking was explained mainly by perceived usefulness, ease of use, trust and perceived risk (Shaikh & Karjaluoto, 2015). In Morocco, for example, mobile-banking use has been linked to technological, social, and organizational conditions (Ziky et al., 2022). Later bibliometric work showed that this research stream expanded unevenly across countries and publication venues (Bahida et al., 2025). A similar movement can be seen in management research, where attention has increasingly shifted from technical feasibility alone toward organizational performance and decision-making (Dwivedi et al., 2021; Mikalef & Gupta, 2021; Ziky et al., 2025).

Auditing approaches the same technological change from a different starting point. Early work considered whether AI could formalize audit procedures, reduce repetitive tasks, or supplement professional effort (Issa et al., 2016). Kokina and Davenport (2017) later documented the emergence of cognitive technologies in audit practice. The issue becomes more demanding once AI influences consequential decisions. At that point, transparency, bias, privacy, reliability, overreliance, and governance are no longer secondary concerns; they can determine whether a system is suitable for use (Kokina et al., 2025). Finance faces a comparable problem. Reviews of explainable AI show that interpretability is especially relevant in credit, fraud, and risk settings, where a model output may need to be justified to a professional reviewer rather than accepted solely because it performs well (Černevičienė & Kabašinskas, 2024; Weber et al., 2024). This requirement is itself contested: some authors consider post hoc explanations of complex models an adequate basis for oversight, whereas others argue that high-stakes decisions call for inherently interpretable models (Barredo Arrieta et al., 2020; Rudin, 2019).

A broad review of AI in banking would now add limited novelty because that landscape is already well mapped. Pattnaik et al. (2024) cover AI and machine learning across financial services; Bouladasse et al. (2026) and Louhab et al. (2026) examine banking applications and performance-related themes; and Kuanova and Otegen (2026) focus more specifically on banking risk management. Auditing and XAI have also been synthesized as distinct domains (Binh, 2025; Černevičienė & Kabašinskas, 2024). The remaining question is narrower. We still know less about the point where AI-enabled financial analytics meet auditing and assurance, risk management, fraud detection, cybersecurity, explainability, and regulatory compliance. That overlap matters because implementation does not end once a model produces a prediction. Financial institutions must still decide how the model will be validated, monitored, documented, challenged, controlled, and, where relevant, audited. The question has also become regulatory: the European regulation on artificial intelligence classifies AI systems used to evaluate the creditworthiness of natural persons as high-risk (Regulation (EU) 2024/1689, 2024), and banking supervisors have begun to examine how banks govern AI and machine-learning models (Basel Committee on Banking Supervision [BCBS], 2022). The problem addressed here therefore lies between technical capability and the assurance requirements attached to AI-supported financial decisions. To our knowledge, this intersection has not yet been the subject of a dedicated bibliometric mapping: existing reviews

treat banking AI, audit AI and XAI as separate objects (see Table 1). This is the gap the present study seeks to fill.

This narrower focus leads to four research questions, formulated as follows. RQ1: How have scientific production and influence at the intersection of AI, auditing and banking risk management evolved over time? RQ2: Which concepts and thematic streams structure this literature? RQ3: On which references and citation relationships does the field draw? RQ4: How are authors and countries connected through scientific collaboration? RQ1 is addressed in Sections 4.1 and 4.2, RQ2 in Section 4.3, RQ3 in Section 4.4 and RQ4 in Section 4.5. Performance analysis and science mapping are combined because these questions concern different dimensions of the literature and cannot be answered by a single indicator.

The study also mobilizes the expression AI-enabled financial risk governance and assurance, whose epistemological status should be clear from the outset. It is a heuristic and interpretive device, used a posteriori to read the patterns revealed by the mapping; it is not a variable or a construct that the bibliometric design operationalizes or measures. It designates the organizational arrangements (validation, monitoring, documentation, control and accountability) through which AI-supported financial decisions can be audited and governed.

The remainder of the article is organized as follows. Section 2 reviews the literature on AI in banking and financial risk management, on AI in auditing and assurance, on explainability, and on the normative frameworks that govern model risk. Section 3 presents the research design, the data source, the screening procedure and the science-mapping techniques. Section 4 reports the results for each research question. Section 5 discusses these results, their theoretical and practical implications, and the limitations of the study, before Section 6 concludes.

2. Literature Review

2.1. Artificial intelligence in banking and financial risk management

Early research on banking digitalization centered on electronic and mobile services, user acceptance, and digitally mediated customer relationships. Moroccan studies illustrate this phase well: perceived usefulness, effort, social influence, facilitating conditions, and crisis-related factors all entered explanations of mobile-banking use (Ziky et al., 2022). Bibliometric evidence subsequently showed that mobile-banking adoption had developed into a recognizable research stream, although publication activity and geographic representation remained uneven (Bahida et al., 2025). AI changes the problem in a more fundamental way. It affects not only the channel through which services are delivered, but also decisions taken inside the institution: customers are classified, risks are estimated, anomalies are flagged, and transactions are screened. Across financial services, the resulting applications include credit assessment, fraud detection, anti-money laundering, underwriting, portfolio management, and operational risk (Pattnaik et al., 2024).

The size of recent AI-banking reviews confirms that the sector-level landscape is already well documented. Louhab et al. (2026), drawing on 114 Scopus papers from 2020 to 2025, identify fraud detection, credit and risk management, digital banking, cybersecurity, customer applications, and explainable AI among the main themes. Bouladasse et al. (2026) analyze a much larger set of 891 Scopus articles covering 2014-2024 and examine trends, intellectual structure, and performance-related topics. Kuanova and Otegen (2026) narrow the lens to banking risk management, using 83 Web of Science articles published between 2020 and 2024. These reviews answer broad questions about where and how AI is studied in banking. They leave a more specific issue less developed: how AI-based risk technologies become connected with audit, assurance, explainability, and the regulatory arrangements that make their organizational use reviewable.

Table 1 compares these previous reviews systematically according to the database used, the period covered, the size of the corpus, the methods employed and the thematic scope. The comparison shows that none of them requires an audit, assurance or control component as an inclusion criterion, which is the distinctive screening rule of the present study.

Table 1: Comparison of the present study with previous reviews

Study	Database	Period	Corpus	Main methods	Thematic scope
Pattnaik et al. (2024)	Scopus	1971-2022	1,045 articles	N-gram and co-occurrence analysis	AI and ML in banking, financial services and insurance
Černevičienė & Kabaši naskas (2024)	Scopus, Web of Science	2005-2022	138 articles	Systematic literature review	XAI applications in finance
Binh (2025)	Scopus, Web of Science	1982-2024	465 publications	Bibliometric analysis, LDA topic modelling	AI and auditing
Bouladasse et al. (2026)	Scopus	2014-2024	891 articles	Performance analysis, intellectual structure	AI in banking, performance-related themes
Louhab et al. (2026)	Scopus	2020-2025	114 documents	Systematic bibliometric review	AI in banking: techniques, applications, challenges
Kuanova & Otegen (2026)	Web of Science	2020-2024	83 articles	Co-occurrence, thematic map, co-citation, coupling	AI in banking risk management
Present study	Scopus	2014-2026	43 documents	Performance analysis, thematic map and evolution, MCA, co- citation, historiography, collaboration	Intersection of AI, audit/assurance and banking risk (explainability, fraud, compliance)

Source: Authors' elaboration based on the cited reviews.

2.2. Artificial intelligence in auditing and assurance

Auditing imposes requirements that go beyond predictive performance. A model may be accurate yet remain difficult to use as audit evidence when its logic is opaque, its behavior is unstable, or the controls around it are weak. That distinction helps explain why early audit research addressed both automation and the formalization of audit procedures (Issa et al., 2016), while Kokina and Davenport (2017) considered how cognitive technologies might alter audit work itself. Since then, robotic process automation, advanced analytics, and machine learning have broadened the toolkit for extracting data, screening transactions, detecting anomalies, and assessing risk. The relevant issue is therefore not simply whether such tools work, but whether their outputs can be incorporated into an auditable process.

The difficulty becomes more pronounced as systems grow more complex. Evidence from professionals in large public accounting firms suggests that relatively simple AI applications are used more widely than complex ones, and the barriers are not only technical. Transparency, explainability, bias, privacy, robustness, reliability, overreliance, and the availability of guidance all become more consequential as model use expands (Kokina et al., 2025). Binh's (2025) review of 465 publications reaches a related conclusion from the literature side, showing growing overlap among AI in auditing, data security, accounting technologies, and machine learning. For assurance practice, this broadens what may need to be examined. Auditors may increasingly have to consider models, data pipelines, and surrounding governance arrangements in addition to transactions and controls.

The literature is not unanimous on these points. One strand treats AI mainly as a means of formalizing audit procedures and supplementing the audit workforce (Issa et al., 2016; Kokina & Davenport, 2017), whereas field evidence points to the risk that professionals over-rely on outputs whose limits they do not fully understand (Kokina et al., 2025). A second divergence concerns the object of assurance. Some contributions focus on auditing with AI, where AI is an audit tool; others address the auditing of AI, where the model itself becomes the object of examination, as in process-based assessments of trustworthy AI (Zicari et al.,

2021). These two perspectives call for different competences and evidence standards, and, as the results below show, they remain only loosely connected in the literature.

2.3. Explainability, trustworthy AI, and regulatory compliance

Explainability becomes especially relevant when financial decisions are both complex and consequential. Lundberg and Lee's (2017) SHAP framework offered one widely adopted way of interpreting model predictions, while Bracke et al. (2019) illustrated how explanation methods could be used in mortgage-default analysis. Later reviews locate many XAI applications in credit, fraud, risk, and financial prediction (Černevičienė & Kabašinskas, 2024; Weber et al., 2024). The appeal is straightforward: greater model complexity may improve predictive performance, yet the same complexity can make it harder for a user, auditor, or regulator to reconstruct why a particular output was produced.

This apparent trade-off between performance and interpretability is itself debated. Proponents of post hoc explanation argue that complex models can be retained provided that techniques such as SHAP make their outputs intelligible (Bracke et al., 2019; Lundberg & Lee, 2017). Rudin (2019) contests this position for high-stakes decisions: explanations of black-box models may be unfaithful to the model they describe, and inherently interpretable models often perform comparably on structured data. Barredo Arrieta et al. (2020) adopt an intermediate position that links explainability to a broader responsible-AI agenda including fairness, accountability and privacy. For auditors and supervisors, the stakes of this debate are concrete, since an explanation that is not faithful to the model cannot serve as reliable audit evidence.

For this reason, explainability increasingly appears in discussions of AI governance rather than as a purely technical feature. Fritz-Morgenthal et al. (2022) connect financial risk management with AI that is explainable, trustworthy, responsible, auditable, and manageable. Giudici et al. (2024) operationalize AI risk through dimensions that include sustainability, accuracy, fairness, and explainability, whereas Zicari et al. (2021) treat trustworthy AI as a broader assessment process. These perspectives do not justify treating explainability as proof of fairness, reliability, or compliance. Its role is narrower: explanation creates a basis for review, challenge, validation, documentation, and accountability. That role is consistent with the present corpus, where explainability repeatedly appears alongside audit, risk, fraud, cybersecurity, and compliance.

2.4. Model-risk governance and regulatory frameworks

The governance perspective discussed above is also anchored in established normative frameworks, to which the bibliometric corpus refers only indirectly. In enterprise risk management, the framework of the Committee of Sponsoring Organizations of the Treadway Commission presents risk management as integral to strategy and performance, with governance and culture, performance monitoring, and information and reporting as core components (Committee of Sponsoring Organizations of the Treadway Commission [COSO], 2017). In banking supervision, the Basel Committee has highlighted the specific challenges raised by AI and machine learning for model risk management: explainability of outcomes, clear accountability involving developers, validators, users and auditors, data governance, and reliance on third-party providers (BCBS, 2022). The European Banking Authority's guidelines on ICT and security risk management require institutions to embed ICT and security risks in their internal control framework, which directly concerns the cybersecurity dimension of AI systems (European Banking Authority [EBA], 2019).

The most explicit framework is the European regulation on artificial intelligence (Regulation (EU) 2024/1689, 2024), known as the EU AI Act, which appears as a retained keyword in the conceptual structure of the corpus (Section 4.3). The regulation adopts a risk-based approach and classifies as high-risk, among others, AI systems intended to evaluate the

creditworthiness of natural persons or establish their credit score, with the explicit exception of systems used to detect financial fraud (Annex III). High-risk systems are subject to requirements on risk management, data governance, technical documentation, record-keeping, transparency towards deployers, human oversight, and accuracy, robustness and cybersecurity (Articles 9 to 15). These requirements translate explainability and auditability into verifiable obligations and provide a normative benchmark against which the bibliometric findings can be interpreted (Section 5.4).

3. Research Methodology

3.1. Research design and data source

We treated the bibliometric exercise as two related tasks rather than a single procedure. The first was descriptive: when were the papers published, where did they appear, who wrote them, and how were citations distributed? The second concerned structure: which concepts, references, and collaborators tend to appear together (Donthu et al., 2021; Zupic & Čater, 2015)? Keeping these tasks separate also helps prevent over-interpretation. A bibliographic link shows an association in the literature; it does not establish that one concept or actor causes another. Most analyses were run in Bibliometrix/Biblioshiny in R (Aria & Cuccurullo, 2017). We moved to R directly when the graphical interface did not give us enough control over temporal slicing or clustering checks.

We chose Scopus as the sole bibliographic source for a practical reason: the topic sits across management, finance, accounting, computer science, and neighboring fields, and Scopus provides structured metadata across all of them. Its export format is also directly compatible with Bibliometrix. Using one database kept the citation and metadata fields internally consistent, although publications indexed only elsewhere are necessarily missed (Mongeon & Paul-Hus, 2016). The magnitude of this coverage bias can be approximated from comparative studies: Singh et al. (2021) report that about 99% of the journals indexed in Web of Science are also indexed in Scopus, so that the omission mainly concerns sources indexed exclusively in other databases rather than the core Web of Science journals. Journal-level overlap does not guarantee document-level overlap, however, and this figure should be read as an order of magnitude rather than an exact estimate. The search was run on 9 September 2026 in TITLE-ABS-KEY. We set no lower publication-year limit and no language restriction; PUBYEAR < 2027 excluded publications after 2026. The substantive query returned 296 records. Restricting document type to articles and reviews left 102 records for duplicate/earlier-version checking and topical screening. The exact query is reproduced below so that the retrieval can be replicated.

TITLE-ABS-KEY (("artificial intelligence" OR "machine learning" OR "deep learning" OR "generative artificial intelligence" OR "generative AI" OR "large language model" OR "natural language processing" OR "neural network*" OR "expert system*" OR "predictive analytics" OR "cognitive computing" OR "robotic process automation") AND (audit* OR auditing OR auditor* OR "financial audit*" OR "internal audit*" OR "external audit*" OR "statutory audit*" OR "digital audit*" OR "continuous audit*" OR "continuous auditing" OR "automated audit*" OR "intelligent audit*" OR "AI-enabled audit*" OR "audit analytics" OR "audit data analytics" OR "audit quality" OR "audit risk" OR "assurance service*" OR "financial assurance") AND (banking OR "banking sector" OR "banking industry" OR "commercial bank*" OR "financial institution*" OR "digital banking" OR "mobile banking" OR "m-banking" OR "online banking" OR "internet banking" OR "electronic banking" OR "e-banking")) AND PUBYEAR < 2027 AND (LIMIT-TO (DOCTYPE , "ar") OR LIMIT-TO (DOCTYPE , "re"))*


3.2. Screening and eligibility

We screened the records in successive passes, using a PRISMA-informed reporting logic

adapted to bibliometric review practice (Page et al., 2021). The TITLE-ABS-KEY search returned 296 records. Limiting the set to articles and reviews reduced this to 102 records, of which 95 were articles and seven reviews. One duplicate/earlier-version indexing record was removed next, leaving 101 unique records. We then read titles, abstracts, and author keywords against the eligibility rules. Fifty-eight records were excluded. The final analytical corpus therefore contains 43 documents: 41 research articles and two reviews. Although the search had no lower year limit, the earliest paper that met all inclusion criteria was published in 2014.

The search string was deliberately broader than the corpus we eventually analyzed. A paper remained in the review only when three elements were substantively connected: an AI-related technique; an auditing, assurance, internal-control, compliance-assurance, auditability, oversight, or closely related control function; and a banking, financial-institution, digital-banking, or financial-services context. This extra screening step was necessary because otherwise the set would have been dominated by prediction studies that merely happened to mention audit-related terms. For traceability, exclusions were coded as follows: wrong or overly broad financial context (E1, $n = 13$); no substantive audit/assurance/control focus (E2, $n = 12$); AI mentioned only incidentally (E3, $n = 3$); prediction-only credit, risk, fraud, or default studies without assurance/control content (E4, $n = 24$); and accounting, finance, or digital-banking studies without substantive assurance/control content (E5, $n = 6$). The resulting corpus should therefore be read as an AI-audit-banking intersection, not as a census of all AI research in banking.

Figure 1: PRISMA-informed bibliometric screening flow



Source: Authors' screening log based on the final Scopus retrieval; PRISMA 2020 reporting logic (Page et al., 2021) adapted to bibliometric screening.

Screening was carried out by the two authors. The first author applied the eligibility rules to all 101 records; the second author then independently re-examined every inclusion and exclusion decision against the same rules and the E1-E5 coding scheme. Divergent or borderline decisions were discussed until consensus was reached, and the final decision and

its rationale were recorded in the screening log. Because this procedure relied on independent verification followed by consensus rather than on parallel blind coding, no inter-rater agreement coefficient (such as Cohen's kappa) was computed; this point is acknowledged among the limitations (Section 5.7). The screening flow is summarized in Figure 1.

3.3. Data preparation and normalization

Before drawing any network, we checked the 43 retained records for missing fields and obvious inconsistencies. They came from 39 publication sources and involved 121 distinct authors. The fields needed for the analyses (title, abstract, source, author, publication year, cited references, and citation counts) were complete. One record had no author keywords (2.33%), and three had no DOI (6.98%). We kept the keyword cleaning conservative. Only clear lexical variants were harmonized: machine-learning/machine learning, cyber security/cybersecurity, risks management/risk management, financial service/financial services, and financial institution/financial institutions. Related terms were not automatically treated as synonyms. Thus, fraud remained separate from fraud detection, explainability from explainable AI, and risk assessment from risk management. This reduced lexical noise without erasing distinctions that later proved useful in the conceptual maps.

3.4. Bibliometric and science-mapping procedures

We kept the main analyses separate instead of forcing them into one network. Performance analysis described annual production, publication sources, author participation, citation patterns, and the most globally cited documents. Conceptual structure was examined with normalized author keywords, thematic mapping, thematic evolution, and MCA. In the thematic map, centrality and density were read as structural properties only. We did not call a low-density theme emerging or declining unless the temporal analysis pointed in the same direction. For thematic evolution, the corpus was split into 2014-2023 and 2024-2026. Because only a small number of themes persisted across both periods, we treated the temporal links as exploratory rather than definitive.

The MCA specification was not fixed in advance. Author keywords had to occur at least twice, which left 17 terms. Bibliometric guidelines do not prescribe a universal frequency cut-off; thresholds are generally adapted to the size of the corpus and to the need to obtain interpretable maps (Donthu et al., 2021; Zupic & Čater, 2015). In a corpus of 43 documents with heterogeneous author keywords, a threshold of two occurrences is the lowest value that excludes idiosyncratic single-document terms. A threshold of three occurrences would, by construction, risk eliminating the low-frequency terms that define the smaller audit, compliance and blockchain streams, and would reduce the map to its technical core. We then compared alternative cluster limits. Five clusters broke the map into very small or singleton groups; two clusters produced the opposite problem by bringing substantively different terms together. Three clusters offered the most defensible compromise between separation and interpretability. A formal resampling test, such as a bootstrap re-estimation of the MCA, was not performed: with 17 terms and 43 documents, resampled subsets would contain too few co-occurrences for the clustering to be re-estimated meaningfully, and a comparison across sub-periods was not feasible because the 2014-2023 period contains only nine documents. Robustness was instead assessed by comparing alternative cluster solutions and by checking that the three-cluster structure was consistent with the thematic map and the thematic evolution, which rely on different algorithms. The first two dimensions accounted for 72.99% of the mapped structure (57.85% for Dimension 1 and 15.14% for Dimension 2). Intellectual structure was examined separately through reference co-citation and historiography. For co-citation, cited papers/references were grouped with Louvain clustering, up to 50 nodes were displayed, isolated nodes were removed, and the minimum edge threshold was set to one. This

produced 39 non-isolated cited-reference nodes ($Q = 0.7683$). We also tested a two-edge threshold. It reduced the network to only two references, which was too restrictive to represent the citation structure of this corpus. Historiography was therefore kept as a separate analysis because it traces direct citation paths rather than shared cited antecedents.

Social structure was assessed through author co-authorship and cross-country collaboration. After isolated authors were removed, 28 authors remained in the co-authorship network; Louvain clustering gave $Q = 0.8442$. Country links came from cross-national co-authorship. We used centrality measures only to describe position in the network: betweenness for brokerage, closeness for relative proximity within connected components, and PageRank for recursive prominence. None was treated as a measure of scientific quality. That distinction is particularly important here because the corpus is small and several components are disconnected; in such a setting, strong-looking network statistics can coexist with limited citation continuity or collaboration.

4. Results

4.1. Descriptive profile and publication dynamics


Publication activity is concentrated near the end of the observation period. The final corpus contains 43 documents published between 2014 and 2026, with seven papers in 2024, 11 in 2025, and 16 in 2026. Those three years alone account for 34 documents, or 79.1% of the corpus. The 2026 count should be read cautiously because the Scopus export was produced before the end of the year. Citation accumulation is also uneven. At the export date, the corpus had received 605 global citations. The mean was 14.1 citations per paper, whereas the median was only three, showing that a relatively small number of papers account for a large share of the total (Table 2 and Figure 2).

Table 2: Descriptive characteristics of the bibliometric corpus

Indicator	Result
Database	Scopus
Documents	43
Publication period	2014-2026
Publication sources	39
Distinct authors	121
Document types	41 articles; 2 reviews
Language	English (43/43)
Global citations at export	605
Mean / median citations per document	14.1 / 3
Documents published in 2024-2026	34 (79.1%)

Source: Authors' calculations from the final Scopus export (N = 43).

Figure 2: Annual scientific production, 2014-2026



Source: Authors' elaboration in Microsoft Word based on the Scopus dataset (annual production computed with Bibliometrix).

4.2. Sources, authors, and influential documents

Publication outlets are widely dispersed. The 43 documents appear in 39 sources, and only four journals contribute two papers each: Computers & Security, Frontiers in Artificial Intelligence, Journal of Risk and Financial Management, and Journal of Theoretical and Applied Information Technology. Every other source appears once. Recurrence is even weaker at author level. Of 121 distinct authors, 120 appear on a single document; Osman M. is the only author represented twice. In this focused corpus, then, neither journals nor authors have yet formed a small core that publishes repeatedly on the topic.

Citation visibility is much more concentrated. Giudici et al. (2024) have 141 global citations and Fritz-Morgenthal et al. (2022) 115; Sánchez-Aguayo et al. (2021) follows with 61. Sachan et al. (2024) and Yaseen and Al-Amarneh (2025) receive 36 and 29 citations, respectively. These five papers together account for 382 of the corpus's 605 global citations (63.1%) (Table 3). That concentration describes visibility beyond the corpus. It should not be confused with influence inside the 43-document set, where local citation links remain sparse.

Table 3: Most globally cited documents in the corpus

Rank	Lead author (year)	Source	Focus / title (shortened)	Global citations
1	Giudici P. (2024)	Expert Systems with Applications	Artificial intelligence risk measurement	141
2	Fritz-Morgenthal S. (2022)	Frontiers in Artificial Intelligence	Financial risk management and explainable, trustworthy, responsible AI	115
3	Sánchez-Aguayo M. (2021)	Computers	Fraud detection, fraud triangle theory and data mining	61
4	Sachan S. (2024)	International Review of Financial Analysis	Human-AI collaboration in financial underwriting	36
5	Yaseen H. (2025)	Journal of Risk and Financial Management	AI-driven fraud detection in banking	29


Source: Authors' calculations from Scopus citation counts at the time of export.

4.3. Conceptual structure

After obvious spelling and lexical variants were harmonized, artificial intelligence and machine learning remained the two main anchors of the author-keyword structure. The terms around them are more informative for the present question: explainability, risk management, fraud detection, cybersecurity, financial risk management, fintech, blockchain, and regulatory compliance recur across the map. In other words, the conceptual space is not simply a catalogue of AI techniques. It links computational methods to financial control, assurance, and governance concerns.

The thematic map sharpens that picture (Figure 3). Machine learning, explainability, and fraud detection occupy a relatively central and developed position. The broader artificial-intelligence-audit-risk-management theme is less internally dense. Fintech and blockchain form a smaller technology-finance theme with comparatively stronger development, while cybersecurity-fraud and explainable-AI-deep-learning sit further from the center. These locations describe structure, not chronology. We therefore avoided calling low-density themes emerging or declining unless the separate temporal analysis supported the same reading.

Figure 3: Thematic map based on normalized author keywords



Source: Authors' calculations based on the Scopus dataset using Bibliometrix/Biblioshiny.

The temporal result points to continuity more than replacement. Machine learning is the retained theme for 2014-2023; in 2024-2026, that earlier stream connects to both artificial intelligence and machine learning. Each link has an inclusion index of 0.50 and a weighted inclusion score of 0.50. The machine-learning-to-artificial-intelligence path shows slightly higher stability and PageRank-based continuity than the machine-learning-to-machine-learning path. Because only a small number of themes recur across the two periods, this should be read as tentative evidence of thematic broadening, not as proof that the field has undergone a definitive shift.

The MCA tells a similar story from another angle (Table 4). Seventeen author keywords occurring at least twice form three clusters, and the first two dimensions account for 72.99% of the mapped structure (57.85% and 15.14%, respectively). The first cluster is much larger than the other two. It combines AI, machine learning, deep learning, explainable AI, explainability, risk management, banking, cybersecurity, fraud, fraud detection, financial risk management, and the EU AI Act. Audit, compliance, and fintech make up the smaller second cluster; blockchain and regulatory compliance form the third. The imbalance matters. Most of the corpus sits in a broad technical-risk core, while organizational assurance and formal compliance remain more specialized streams.

Table 4: Conceptual clusters identified by multiple correspondence analysis

Cluster	Interpretive label	Main retained terms	Interpretation
1	AI-enabled risk, fraud and explainable financial analytics	Artificial intelligence; machine learning; deep learning; explainable AI; explainability; risk management; banking; cybersecurity; fraud; fraud detection; financial risk management; EU AI Act	Dominant conceptual core linking AI techniques with risk, fraud, banking and explainability/governance concerns.
2	Audit, compliance and digital financial transformation	Audit; compliance; fintech	Specialized stream linking assurance/control functions with digital financial transformation.
3	Blockchain and regulatory compliance infrastructure	Blockchain; regulatory compliance	Small specialized stream focused on digital infrastructure and formal compliance.


Source: Authors' MCA based on normalized author keywords; 17 retained terms; Dim. 1 = 57.85%, Dim. 2 = 15.14%.

4.4. Intellectual structure

The co-citation result looks highly modular on paper, but the graph itself is less settled (Figure 4). After isolated nodes were removed, 39 cited-reference nodes remained and Louvain clustering yielded $Q = 0.7683$. Several components are nevertheless small and disconnected. The value therefore captures both thematic differentiation and fragmentation. It would be too strong to read it as evidence that mature intellectual schools have already formed.

A small number of references do most of the bridging between otherwise separated parts of the graph. Bracke et al. (2019) has the highest visible betweenness centrality (78), followed by Zicari et al. (2021) (70). Three further high-betweenness nodes cannot be attributed to a precise bibliographic record because the cited-reference strings exported by Scopus are incomplete: a node labelled “Chen Z.” (60), exported without publication year or title; a generic node labelled “artificial intelligence (AI)” (50), which corresponds to a truncated reference string rather than to an identifiable author; and a 2021 reference on AML/CFT technologies (48) whose author field could not be resolved unambiguously. These three elements are therefore reported as unresolved node labels and are not listed among the references. Lundberg and Lee (2017) also link explainability-oriented work with the wider network. At community level, four substantive streams can be distinguished: AI, explainability and regulatory governance; XAI and AML/CFT technologies; fraud risk, internal control and banking risk management; and automation, trustworthy AI and digital control. Some cited-reference strings are incomplete or inconsistently standardized, so the community-level pattern is more dependable than a literal reading of every individual node label. To quantify this limitation, at least six of the 39 nodes (15.4%) carry incomplete or ambiguous labels: the three nodes mentioned above and the labels “Huang F.”, “Arrieta A.B.” and “Regulation (EU)”, displayed without publication year. The last two most probably correspond to Barredo Arrieta et al. (2020) and to a European regulation, but the exported strings do not allow a definitive match.

Figure 4: Reference co-citation network (39 non-isolated cited-reference nodes; $Q = 0.7683$)



Source: Authors' calculations based on the Scopus dataset using Bibliometrix/Biblioshiny.

The historiograph is much thinner. Only four documents form identifiable direct-citation paths: Fritz-Morgenthal et al. (2022) links to Jagadhabhi (2025), while Tanbour et al. (2025) links to Johri et al. (2026). These short chains indicate limited citation continuity inside the retained corpus. That does not contradict the co-citation map. The two analyses capture


different relationships: papers can rely on the same intellectual antecedents even when long direct-citation chains have not yet developed among them.

4.5. Social structure

The co-authorship network is similarly fragmented. It contains 28 non-isolated authors and has a modularity of $Q = 0.8442$, yet the visual structure is dominated by small, disconnected teams. The authors shown in the centrality table all have betweenness values of zero. Here, high modularity mainly reflects separation among teams rather than integration into a single collaborative community.

International collaboration is present, but repetition is limited. We identify 22 country-to-country links. Jordan connects with Iraq, Kazakhstan, Malaysia, Saudi Arabia, Thailand, the United Arab Emirates, and the United Kingdom; Saudi Arabia, the United Kingdom, the United Arab Emirates, India, and Malaysia also recur across cross-national pairs. Most bilateral links appear only once. The world map (Figure 5) therefore needs a careful reading: country shading represents publication activity, whereas connecting lines show international co-authorship.

Figure 5: Countries' collaboration world map



Source: Authors' calculations based on the Scopus dataset using Bibliometrix/Biblioshiny.

4.6. Integrated synthesis

Taken together, the four analytical dimensions do not move at the same pace. Conceptually, recent papers increasingly place AI and machine learning alongside risk, fraud, explainability, cybersecurity, audit, and compliance. Citation continuity and collaboration are less developed. The literature therefore looks more integrated at the level of themes than at the level of direct citation lineages or stable research partnerships.

5. Discussion

5.1. Recent expansion and interdisciplinary positioning

The post-2023 concentration of papers may look like rapid growth, but the evidence does not support that stronger claim. The corpus contains only 43 documents, and the 2026 count is incomplete. A narrower conclusion is safer: this particular intersection has become more visible in recent years. That reading is consistent with the wider expansion of AI research in

financial services (Pattnaik et al., 2024) and recent bibliometric work on AI in banking (Bouladasse et al., 2026; Louhab et al., 2026). The difference here is the screening boundary. Prediction-focused and general digital-banking studies were excluded whenever audit, assurance, internal control, compliance assurance, or auditability remained peripheral. The final $N = 43$ therefore reflects the research question rather than weak retrieval.

The spread across publication outlets also helps explain the field's interdisciplinary character. Forty-three papers are distributed across 39 sources, ranging from accounting and auditing journals to outlets in AI, cybersecurity, financial risk, information systems, law, and fintech. The dispersion is substantive as well as bibliographic. The same AI system can be studied as a technical artifact, a decision aid, an object of internal control, and a regulatory concern. That breadth is useful, but it makes cumulative knowledge harder to build because similar questions about oversight and accountability may develop in disciplinary conversations that rarely cite or collaborate with one another.

5.2. From AI application to AI-enabled financial risk governance

The central position of AI and machine learning is not surprising; both concepts are built into the retrieval strategy. The vocabulary around them is more revealing. Fraud detection, financial risk management, cybersecurity, explainability, and regulatory compliance repeatedly sit close to the technical core. Broad banking reviews already document many of these applications (Kuanova & Otegen, 2026; Louhab et al., 2026). In this corpus, however, they are tied more directly to audit and assurance. The substantive question therefore shifts from whether AI can be used in banking to how AI-supported decisions are validated, monitored, challenged, and incorporated into an assurance process.

We use the expression AI-enabled financial risk governance and assurance as a shorthand for that convergence; as stated in the introduction, it is an interpretive device and not a separate bibliometric construct that we measured. The phrase captures the meeting point among model-based risk applications, compliance, explainability, auditability, and oversight. This reading is consistent with Fritz-Morgenthal et al. (2022) and Giudici et al. (2024), both of whom place AI use within broader risk and governance considerations. The thematic-evolution result adds some support, although only cautiously. In the later period, the earlier machine-learning stream broadens toward an artificial-intelligence label. One plausible reading is that attention is moving from individual algorithms toward the organizational systems in which those algorithms are used.

5.3. Explainability as an interface between prediction and accountability

Explainability occupies one of the clearest connecting positions in the conceptual structure. It appears near machine learning, fraud detection, financial risk, and regulatory compliance. In this literature, XAI seems to do more than solve a technical problem: it creates a route through which model outputs can be opened to professional review. That reading is consistent with reviews locating XAI mainly in credit, risk, and fraud applications (Černevičienė & Kabašinskas, 2024; Weber et al., 2024) and with the continuing influence of SHAP (Lundberg & Lee, 2017). Bracke et al. (2019) provide a concrete financial example through explanation methods applied to default-risk decisions.

For auditors and risk professionals, explanation solves only part of the problem. A model can be accurate and still be difficult to rely on if it cannot be validated, behaves inconsistently across populations, depends on problematic data, or conflicts with regulatory expectations. Explanations can make review and challenge easier, but they do not establish fairness, trustworthiness, or compliance. Rudin's (2019) warning that post hoc explanations may not be faithful to the model they describe reinforces this caution. Zicari et al. (2021) make a similar point in their process-based treatment of trustworthy AI, where responsible assessment

depends on several dimensions at once. The network therefore supports a limited conclusion: explainability can enable scrutiny, but responsible AI requires more than explanation alone.

5.4. Auditing, assurance capability, and converging risk categories

The MCA exposes a separation that is easy to overlook when the themes are read only as a list. Technical and risk-oriented terms dominate the largest cluster, whereas audit, compliance, and fintech form a smaller organizational stream. That separation raises a practical distinction between the capacity to deploy AI and the capacity to assure its use. We use AI capability for the former, that is, prediction, classification, automation, anomaly detection, and decision support. Assurance capability refers to the surrounding organizational practices: validation, ongoing monitoring, audit trails, control testing, documentation, and clear accountability for outputs. The distinction is interpretive rather than empirically validated, but it helps explain why audit and compliance remain partly outside the technical core. Table 5 synthesizes the dimensions that distinguish the two capabilities.

Table 5: AI capability versus assurance capability

Dimension	AI capability	Assurance capability
Purpose	Produce predictions, classifications, scores and alerts	Provide reasonable assurance that AI outputs are reliable, explainable and compliant
Core activities	Model development and training, anomaly detection, scoring, automation, decision support	Independent validation, ongoing monitoring, control testing, audit trails, documentation
Key actors	Data scientists, business lines, technology providers	Risk management, compliance, internal audit, external auditors, supervisors
Evaluation criteria	Accuracy, recall, speed, cost reduction	Explainability, traceability, robustness, fairness, accountability
Risks addressed	Credit, fraud, cyber and operational risks	Model risk, governance risk, regulatory and conduct risks
Related clusters	MCA Cluster 1 (AI, risk, fraud, explainability)	Cluster 2 (audit, compliance, fintech) and Cluster 3 (blockchain, regulatory compliance)
Normative anchors	Technical performance standards	COSO (2017); BCBS (2022); EBA (2019); EU AI Act, Articles 9-15 and 26

Source: Authors' elaboration based on the MCA results (Table 4) and the literature reviewed in Section 2.

Evidence from audit practice makes this interpretation plausible. Kokina et al. (2025) show that widespread use of relatively simple AI does not imply equally broad adoption of complex systems. As model use becomes more consequential, explainability, robustness, privacy, bias, and overreliance become harder to set aside. Earlier contributions emphasized AI as a way to formalize and automate audit work (Issa et al., 2016; Kokina & Davenport, 2017). The more recent issue is broader: systems adopted by the client organization may themselves become objects of assurance. AI therefore changes both the auditor's toolkit and the boundary of what may need to be audited.

The regulatory implications of this distinction become concrete when the findings are read against the EU AI Act. For credit-scoring systems classified as high-risk, providers must establish a risk-management system, prepare technical documentation, ensure automatic logging of events, supply deployers with instructions that allow them to interpret outputs, design the system for effective human oversight, and achieve appropriate accuracy, robustness and cybersecurity (Regulation (EU) 2024/1689, 2024, Articles 9-15). Deployers, including banks, must use such systems in accordance with these instructions, assign human oversight to competent persons and keep the logs generated (Article 26); those that assess creditworthiness must also carry out a fundamental-rights impact assessment (Article 27). Each of these obligations belongs to assurance capability rather than to AI capability, and

each can be mapped onto the audit practices described in Section 2.2: technical documentation and logs become audit evidence, human oversight becomes a control to be tested, and robustness requirements call for independent validation. Conversely, fraud-detection systems, a prominent application in Cluster 1, are excluded from the high-risk category of creditworthiness assessment, so that their assurance rests more heavily on prudential and internal-control frameworks (BCBS, 2022; EBA, 2019) than on the AI Act itself.

Fraud, cybersecurity, financial risk, and compliance also sit close together for practical reasons. These risks rarely stay in separate operational silos. A cyber weakness can facilitate fraud; suspicious transactions may trigger AML/CFT obligations; and the models used to detect either problem create model and governance risks of their own. Sachan et al. (2024) add the human dimension by showing why decision consistency and human-AI collaboration matter in lending. Read together, these links suggest that technology-enabled assurance may increasingly need to examine integrated control environments rather than each model or risk category in isolation.

5.5. Intellectual and collaborative fragmentation

The intellectual network calls for the same caution. $Q = 0.7683$ indicates separation among reference communities, but the graph is visibly fragmented and some cited-reference labels remain imperfectly standardized (at least six of the 39 nodes, as quantified in Section 4.4). Calling these communities mature intellectual schools would therefore go too far. More informative are the references that connect otherwise separate conversations, particularly work on explainability, trustworthy AI, AML/CFT, and risk governance. Historiography reaches a similar conclusion from another angle. Co-citation shows shared antecedents, while only four documents form direct citation paths within the corpus. A young interdisciplinary field can draw on common parent literatures well before it develops long internal citation chains of its own.

Collaboration is fragmented as well. The author network has a high modularity value ($Q = 0.8442$), yet the graph is made up mainly of small teams, and most of the 22 international country pairs occur only once. This matters for AI-enabled assurance because the underlying questions cross disciplinary boundaries. Model validation, audit evidence, internal control, legal accountability, cybersecurity, and implementation are difficult to address within a single research tradition. More sustained collaboration among accounting, finance, computer science, cybersecurity, and regulatory scholars would therefore add more to the field than simply increasing the number of isolated application studies.

5.6. Contributions and implications

The contribution becomes clearer when the findings are read together rather than as a sequence of bibliometric outputs. By requiring a substantive audit and assurance component during screening, the study isolates a narrower intersection than broad AI-banking bibliometries or banking-risk reviews (Bouladasse et al., 2026; Kuanova & Otegen, 2026; Louhab et al., 2026). Within that intersection, conceptual themes are converging faster than citation lineages and collaboration networks. That mismatch provides the empirical basis for two interpretive propositions developed in the discussion: the distinction between AI capability and assurance capability, and the broader idea of AI-enabled financial risk governance and assurance.

These propositions have a theoretical scope beyond the description of the corpus. They suggest that the adoption of AI in banking cannot be explained by technical performance alone, and that the reliance placed on AI outputs by auditors and supervisors is likely to depend on the assurance capability of the institution rather than on the predictive accuracy of

its models. Formulated in this way, the distinction becomes testable and connects the bibliometric evidence with the normative frameworks reviewed in Section 2.4.

For practice, the implication is straightforward: governance cannot simply be added after deployment. Banks using AI for consequential decisions need processes for model validation, explanation, data quality, cyber and fraud monitoring, and internal control. Auditors face a related expansion of scope because models, automated controls, data pipelines, and AI-supported decisions may enter the assurance perimeter alongside transactions and financial statements. Regulators encounter the same convergence at another level. Explainability, human oversight, audit trails, model risk, and accountability may appear as separate requirements on paper, but they interact inside the same operational system.

5.7. Limitations and future research

The conclusions remain bounded by the design. The review relies on Scopus alone and deliberately narrows the final corpus to studies with a substantive audit, assurance, or control component. The 2026 publication count is incomplete, and recent papers have had less time to accumulate citations. Network results also depend on metadata cleaning and threshold choices; some cited-reference strings are incomplete or heterogeneous, and disconnected components partly account for the high modularity values. Screening relied on independent verification followed by consensus rather than on parallel blind coding, so that no inter-rater agreement coefficient could be reported, and the stability of the MCA clusters was assessed by comparing alternative solutions rather than by resampling. More fundamentally, co-occurrence and co-citation describe relationships within a body of literature. They do not show that AI improves audit quality, reduces financial risk, or enhances organizational performance.

The next step should therefore be empirical rather than bibliometric. Future work could test how different explanation formats affect auditors' reliance on AI-generated evidence, compare how regulatory regimes shape model governance in banks, or examine when professionals appropriately accept, question, or override algorithmic recommendations. Studies that consider financial, fraud, cyber, operational, and compliance risk together would also extend the control perspective developed here. Longitudinal research, experiments, fieldwork, mixed methods, and cross-country comparisons can provide evidence that bibliometric mapping cannot. Emerging economies, including African settings, deserve particular attention because institutional capacity and regulatory arrangements may alter how assurance around AI is designed and practiced. Morocco offers a particularly relevant setting in this respect. Earlier work has documented the determinants of mobile-banking use among Moroccan customers (Ziky et al., 2022) and the development of the mobile-banking literature (Bahida et al., 2025). Extending this research programme from adoption to the assurance of AI-based systems would make it possible to observe how Moroccan banks, their auditors and the supervisory authority build assurance capability in a context where AI-specific regulation is still emerging.

6. Conclusion

This review focuses on 43 Scopus-indexed documents at the intersection of AI, auditing, banking, financial risk management, fraud detection, explainability, cybersecurity, and regulatory compliance. The literature is recent and spread across many journals and authors, but a recognizable conceptual structure has begun to form.

The results answer the four research questions as follows. Regarding RQ1, production is recent and accelerating, with 79.1% of the documents published between 2024 and 2026, while citation influence is concentrated in five papers that account for 63.1% of global citations. Regarding RQ2, the field is organized around a dominant AI-risk-fraud-

explainability core and two smaller streams devoted to audit, compliance and fintech, and to blockchain and regulatory compliance. Regarding RQ3, the field draws on shared antecedents in explainability, trustworthy AI, AML/CFT and risk governance, but co-citation communities are fragmented and direct citation chains are short. Regarding RQ4, collaboration takes place in small, disconnected teams, and most international links occur only once. The field is therefore more integrated conceptually than it is intellectually or socially.

The more consequential result is not simply what these systems can do, but what organizations need before they can rely on them. Across the mapped literature, an implicit distinction appears between technical capability and assurance requirements. Prediction, classification, automation, anomaly detection, and decision support sit on one side; validation, monitoring, documentation, control testing, audit trails, and accountability sit on the other. We describe this distinction as AI capability versus assurance capability. Explainability can help connect the two by making professional scrutiny easier, but explanation alone does not establish accuracy, fairness, trustworthiness, or compliance.

These results translate into concrete implications for each category of actor. Banks could maintain an inventory of the AI models used in consequential decisions, with a designated owner, documented validation results and explicit criteria for human override for each model. Internal and external auditors could integrate items on AI-model governance (completeness of the model inventory, independence of validation, retention of logs, traceability of overrides) into their existing audit risk matrices. Supervisors, including Bank Al-Maghrib in the Moroccan context, could ask institutions to report their use of high-impact AI models and to demonstrate that the explanation techniques they rely on are faithful to the underlying models. AI providers, finally, could supply the documentation and logs that banks need to meet their own assurance obligations.

These conclusions remain bounded by a deliberately narrow, single-database design whose limitations are detailed in Section 5.7. Building on the authors' earlier work on mobile banking in Morocco, a natural extension is to examine empirically how Moroccan banks and their auditors develop assurance capability for AI-based systems. The question that follows from this review is not simply whether AI improves prediction or automation, but whether AI-supported financial systems can be examined, challenged, governed, secured, and held accountable in actual organizational practice.

Data Availability Statement. The bibliographic data analyzed in this study were exported from Scopus and remain subject to the database provider's licensing and copyright conditions. Derived analytical data, the screening log, keyword-normalization thesaurus, study-level coding matrices, and relevant R scripts may be made available by the corresponding author upon reasonable request. The structured extraction workbook and coding matrices are available as supplementary files for peer-review purposes.

References

- (1). 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>
- (2). Bahida, H., Ziky, R., Chakra, R., Oulamine, A., & Massiki, A. (2025). The adoption of mobile banking services: A comprehensive bibliometric analysis. *International Journal of Innovative Research and Scientific Studies*, 8(5), 179-191. <https://doi.org/10.53894/ijirss.v8i5.8595>
- (3). Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F.

- (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- (4). Basel Committee on Banking Supervision. (2022). *Newsletter on artificial intelligence and machine learning*. Bank for International Settlements. https://www.bis.org/publ/bcbs_nl27.htm
- (5). Binh, N. T. T. (2025). Transforming auditing in the AI era: A comprehensive review. *Information*, 16(5), 400. <https://doi.org/10.3390/info16050400>
- (6). Bouladasse, H., El Ganich, S., & Yahyaoui, T. (2026). Artificial intelligence in banking: A bibliometric analysis of research trends, intellectual structure, and performance-related themes. *Journal of Risk and Financial Management*, 19(8), 581. <https://doi.org/10.3390/jrfm19080581>
- (7). Bracke, P., Datta, A., Jung, C., & Sen, S. (2019). *Machine learning explainability in finance: An application to default risk analysis* (Staff Working Paper No. 816). Bank of England. <https://doi.org/10.2139/ssrn.3435104>
- (8). Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. *Artificial Intelligence Review*, 57, 216. <https://doi.org/10.1007/s10462-024-10854-8>
- (9). Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*. COSO.
- (10). 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>
- (11). Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., ... Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
- (12). European Banking Authority. (2019). *Guidelines on ICT and security risk management* (EBA/GL/2019/04). European Banking Authority.
- (13). Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. *Frontiers in Artificial Intelligence*, 5, 779799. <https://doi.org/10.3389/frai.2022.779799>
- (14). Giudici, P., Centurelli, M., & Turchetta, S. (2024). Artificial intelligence risk measurement. *Expert Systems with Applications*, 235, 121220. <https://doi.org/10.1016/j.eswa.2023.121220>
- (15). Issa, H., Sun, T., & Vasarhelyi, M. A. (2016). Research ideas for artificial intelligence in auditing: The formalization of audit and workforce supplementation. *Journal of Emerging Technologies in Accounting*, 13(2), 1-20. <https://doi.org/10.2308/jeta-10511>
- (16). Jagadhabi, N. (2025). Explainable AI-driven decision support for strategic risk management in financial services. *International Academic Journal of Science and Engineering*, 12(3), 125-144. <https://doi.org/10.71086/iajse/v12i3/iajse1254>
- (17). Johri, A., Sayal, A., Chong, K. M., Khoja, M., N, C., Jha, J., & Tyagi, N. (2026). Enhancing audit quality and reducing costs: The impact of AI in banking and financial services. *Frontiers in Artificial Intelligence*, 8, 1718854. <https://doi.org/10.3389/frai.2025.1718854>
- (18). Kokina, J., Blanchette, S., Davenport, T. H., & Pachamanova, D. (2025). Challenges and opportunities for artificial intelligence in auditing: Evidence from the field.

- International Journal of Accounting Information Systems*, 56, 100734. <https://doi.org/10.1016/j.accinf.2025.100734>
- (19). Kokina, J., & Davenport, T. H. (2017). The emergence of artificial intelligence: How automation is changing auditing. *Journal of Emerging Technologies in Accounting*, 14(1), 115-122. <https://doi.org/10.2308/jeta-51730>
 - (20). Königstorfer, F., & Thalmann, S. (2020). Applications of artificial intelligence in commercial banks: A research agenda for behavioral finance. *Journal of Behavioral and Experimental Finance*, 27, 100352. <https://doi.org/10.1016/j.jbef.2020.100352>
 - (21). Kuanova, L. A., & Otegen, A. N. (2026). Artificial intelligence in banking risk management: A bibliometric analysis. *International Journal of Financial Studies*, 14(4), 93. <https://doi.org/10.3390/ijfs14040093>
 - (22). Louhab, A., Yaacoubi, A., & Azouazi, M. (2026). Mapping the evolution of artificial intelligence in banking: A systematic bibliometric review of research trends, techniques, applications, and emerging challenges. *International Journal of Accounting, Finance, Auditing, Management and Economics*, 7(8), 415-437. <https://ijafame.org/index.php/ijafame/article/view/2523>
 - (23). Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
 - (24). Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. *Information & Management*, 58(3), 103434. <https://doi.org/10.1016/j.im.2021.103434>
 - (25). Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106, 213-228. <https://doi.org/10.1007/s11192-015-1765-5>
 - (26). Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
 - (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). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). (2024). *Official Journal of the European Union*, L series, 12 July 2024. <http://data.europa.eu/eli/reg/2024/1689/oj>
 - (29). Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
 - (30). Sachan, S., Almaghrabi, F., Yang, J.-B., & Xu, D.-L. (2024). Human-AI collaboration to mitigate decision noise in financial underwriting: A study on FinTech innovation in a lending firm. *International Review of Financial Analysis*, 93, 103149. <https://doi.org/10.1016/j.irfa.2024.103149>
 - (31). Sánchez-Aguayo, M., Urquiza-Aguilar, L., & Estrada-Jiménez, J. (2021). Fraud detection using the fraud triangle theory and data mining techniques: A literature review. *Computers*, 10(10), 121. <https://doi.org/10.3390/computers10100121>
 - (32). Shaikh, A. A., & Karjaluoto, H. (2015). Mobile banking adoption: A literature review. *Telematics and Informatics*, 32(1), 129-142. <https://doi.org/10.1016/j.tele.2014.05.003>

- (33). Singh, V. K., Singh, P., Karmakar, M., Leta, J., & Mayr, P. (2021). The journal coverage of Web of Science, Scopus and Dimensions: A comparative analysis. *Scientometrics*, 126(6), 5113-5142. <https://doi.org/10.1007/s11192-021-03948-5>
- (34). Tanbour, K. M., Ben Saada, M., Nour, A. I., & Elnaas, N. K. (2025). Integrating artificial intelligence into risk management frameworks: A mixed-methods analysis of the Palestinian banking sector. *Journal of Financial Reporting and Accounting*. Advance online publication. <https://doi.org/10.1108/JFRA-06-2025-0458>
- (35). Weber, P., Carl, K. V., & Hinz, O. (2024). Applications of explainable artificial intelligence in finance: A systematic review of finance, information systems, and computer science literature. *Management Review Quarterly*, 74, 867-907. <https://doi.org/10.1007/s11301-023-00320-0>
- (36). Yaseen, H., & Al-Amarneh, A. (2025). Adoption of artificial intelligence-driven fraud detection in banking: The role of trust, transparency, and fairness perception in financial institutions in the United Arab Emirates and Qatar. *Journal of Risk and Financial Management*, 18(4), 217. <https://doi.org/10.3390/jrfm18040217>
- (37). Zicari, R. V., Brodersen, J., Brusseau, J., Düdler, B., Eichhorn, T., Ivanov, T., Kararigas, G., Kringen, P., McCullough, M., Möslin, F., Mushtaq, N., Roig, G., Stürtz, N., Tolle, K., Tithi, J. J., van Halem, I., & Westerlund, M. (2021). Z-Inspection®: A process to assess trustworthy AI. *IEEE Transactions on Technology and Society*, 2(2), 83-97. <https://doi.org/10.1109/TTS.2021.3066209>
- (38). Ziky, R., Abriane, A., Bahida, H., Charef, F., & Bouzahir, B. (2022). Les déterminants de l'utilisation du mobile banque par les clients marocains : Cas des clients du Souss Massa. *International Journal of Accounting, Finance, Auditing, Management and Economics*, 3(1-2), 503-521. <https://doi.org/10.5281/zenodo.5910159>
- (39). Ziky, R., Bahida, H., & Abriane, A. (2025). The impact of artificial intelligence on business performance: A bibliometric analysis of publication trends. *Moroccan Journal of Quantitative & Qualitative Research*, 7(1). <https://doi.org/10.48379/IMIST.PRSM/mjqr-v7i1.53210>
- (40). Zupic, I., & Cater, T. (2015). Bibliometric methods in management and organization. *Organizational Research Methods*, 18(3), 429-472. <https://doi.org/10.1177/1094428114562629>

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process:

Artificial intelligence tools were used exclusively for translation and language editing purposes, without contributing to the research design, data analysis, or interpretation of results. The authors take full responsibility for the final content of the article and have reviewed and verified all modifications made.